

Optimal Transport Guided Correlation Assignment for Multimodal Entity Linking

Zefeng Zhang*, Jiawei Sheng*, Chuang Zhang†,

Yunzhi Liang, Wenyuan Zhang, Siqi Wang, Tingwen Liu

Institute of Information Engineering, Chinese Academy of Sciences. Beijing, China
School of Cyber Security, University of Chinese Academy of Sciences. Beijing, China
{zhangzefeng, shengjiawei, zhangchuang}@iie.ac.cn
{liangyunzhi, zhangwenyuan, wangsiqi2022, liutingwen}@iie.ac.cn

Abstract

Multimodal Entity Linking (MEL) aims to link ambiguous mentions in multimodal contexts to entities in a multimodal knowledge graph. A pivotal challenge is to fully leverage multi-element correlations between mentions and entities to bridge modality gap and enable fine-grained semantic matching. Existing methods attempt several local correlative mechanisms, relying heavily on the automatically learned attention weights, which may over-concentrate on partial correlations. To mitigate this issue, we formulate the correlation assignment problem as an optimal transport (OT) problem, and propose a novel MEL framework, namely OT-MEL, with OT-guided correlation assignment. Thereby, we exploit the correlation between multimodal features to enhance multimodal fusion, and the correlation between mentions and entities to enhance fine-grained matching. To accelerate model prediction, we further leverage knowledge distillation to transfer OT assignment knowledge to attention mechanism. Experimental results show that our model significantly outperforms previous state-of-the-art baselines and confirm the effectiveness of the OT-guided correlation assignment.¹

1 Introduction

Entity Linking (EL) is an important yet challenging task in knowledge acquisition, which can facilitate applications such as information retrieval (Liu et al., 2018), question answering (Xiong et al., 2019) and dialogue systems (Jiang et al., 2023). In recent years, with the rapid increase of multimodal content, Multimodal Entity Linking (MEL) has been proposed, which aims to link ambiguous mentions in multimodal contexts to entities in the multimodal knowledge graph (KG). As shown in

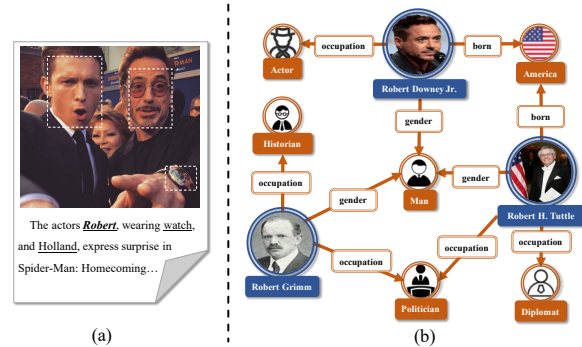


Figure 1: Examples of MEL: (a) the given mention “Robert” in a multimodal content, and (b) the candidate entities in the multimodal KG.

Figure 1, given a mention Robert in texts with related images, the MEL model requires the mention to be corresponded to the entity Robert_Downey in the multimodal KG.

Since the mention can be ambiguous with limited textual information, it is crucial to enrich the text with its related visual images. To utilize multimodal information for MEL, a pivotal challenge is to leverage *multi-element correlations* within entity and mention contexts: 1) *the multi-element correlations between modalities*, where the textual tokens and visual patches may have correlative correspondence. As shown in Figure 1, the token Robert and Holland correspond to different local areas of human faces in the image. These correlations serve as message-passing pivots between the text and image, which can be valuable to bridge the semantic gap of modalities (Luo et al., 2023). 2) *the multi-element correlations between mentions and entities*, where the context of mentions and entities can have multiple correlative cues for semantic matching. As shown in Figure 1, the mentions and entity both involve token Robert, and both the mention and the entity involve the word Actor. These correlations depict similarities between the mention and entity, providing explicit evidence for EL.

* These authors contribute equally to this work.

† Corresponding Authors.

¹Source code of this paper could be obtained from <https://github.com/zhangzefeng/OT-MEL>

To facilitate these multi-element correlations, existing studies propose sophisticated correlative mechanisms, such as feature concatenation (Adjali et al., 2020), attention-based modules (Moon et al., 2018; Wang et al., 2022a), hierarchical gates (Luo et al., 2023; Wang et al., 2022a) and graph-based aggregation (Xing et al., 2023). However, these models learn the correlations of elements locally, and rely heavily on the automatically learned correlation assignments. This design may over-concentrate on part of elements, dominating the correlations without global consideration. For example, all entities in Figure 1 have surface name Robert the same to mention Robert, which may attract the correlation assignment and conceal crucial correlations of token Actor in texts and the visual patch of human face in images.

To mitigate the above issues, we formulate the multi-element correlation assignment as an *optimal transport (OT)* problem. Taking the multimodal correlation as an example, given the cost of each pair of textual tokens and visual patches, the OT problem aims to find an optimal transport plan to transfer the token to a patch with assigned weights, such that the total cost is minimal. In contrast to the automatically learned attention map, the OT plan globally considers the correlation between source and target elements, surpassing the degraded dominant correspondence (Liu et al., 2020a).

Following the above idea, we propose a novel MEL framework with OT guided correlation assignment, termed **OT-MEL**. Specifically, we first adopt multimodal feature encoders to represent textual tokens and visual patches for both mentions and entities. Afterward, we propose an OT-based correlation assignment method between textual tokens and visual patches, with which we further integrate multimodal features and bridge the semantic gap of modalities. Additionally, we also build unimodal correlations between mentions and entities with OT-based correlation assignment. By measuring the similarity of the multimodal and unimodal features between mentions and entities, the model derives the overall matching score. To enhance efficiency, we further propose distillation between OT assignment to attention map, such that we can use attention to approximate OT plan. The contributions of this paper are three-fold:

- We introduce optimal transport to capture complicated correlations in MEL. To our knowledge, we are the first to guide MEL with OT.

- We propose OT-MEL, a novel framework to consider the multimodal and unimodal correlations, which can be further extended by knowledge distillation from OT plan to attention mechanism.
- Empirical studies indicate that our model outperforms previous state-of-the-art methods on three widely-used benchmarks and extensive analysis confirm the effectiveness of our model.

2 Preliminary

2.1 Task Formulation

The task of MEL aims to correspond a given mention from text within a multimodal context to its true entity in a multimodal KG. Generally, the KG involves entities $e \in \mathcal{E}$, where each entity $e \triangleq \{e^n, e^v, e^a\}$ can be associated with entity name e^n , visual image e^v , and a textual attribute set e^a ². Besides, the given mention from texts can be summarized as $m \triangleq \{m^n, m^c, m^v\}$, which is associated with mention word m^n , context sentence m^c , and visual image m^v , respectively.

To achieve this task, for each given mention, the model requires to identify the true entity from entity candidates, which can be formulated as:

$$\theta^* = \arg \max_{\theta} \sum_{(m,e) \in \mathcal{D}} \log p_{\theta}(e|m, \mathcal{E}), \quad (1)$$

where θ^* denotes the final model. Following previous studies (Luo et al., 2023), we use all entities in KG as the candidates.

2.2 Optimal Transport

Optimal Transport (OT) aims to find a transport plan with minimal total cost, which transports the density distribution of a group of elements to that of another group, given the cost of each element pair (Villani et al., 2009). Formally, given element x_i in one group, y_j in another group and C_{ij} the cost between them, the problem can be formulated as the following linear programming problem:

$$\begin{aligned} \mathbf{\Pi}^* &= \arg \min_{\mathbf{\Pi} \in \mathbb{R}_+^{n \times m}} \sum_{i=1}^m \sum_{j=1}^n C_{ij} \Pi_{ij}, \\ \text{s.t. } \mathbf{\Pi} \mathbf{1}_n &= \boldsymbol{\mu}, \quad \mathbf{\Pi}^\top \mathbf{1}_m = \boldsymbol{\nu}, \end{aligned} \quad (2)$$

where $\mathbf{\Pi}^*$ is the optimal transport plan, where Π_{ij} is the amount of information from x_i to y_j . $\mathbf{1}_n$

²Following most existing MEL studies (Wang et al., 2022a; Luo et al., 2023), we also treat the attributes as textual information of the entities due to the large-scale KG consideration.

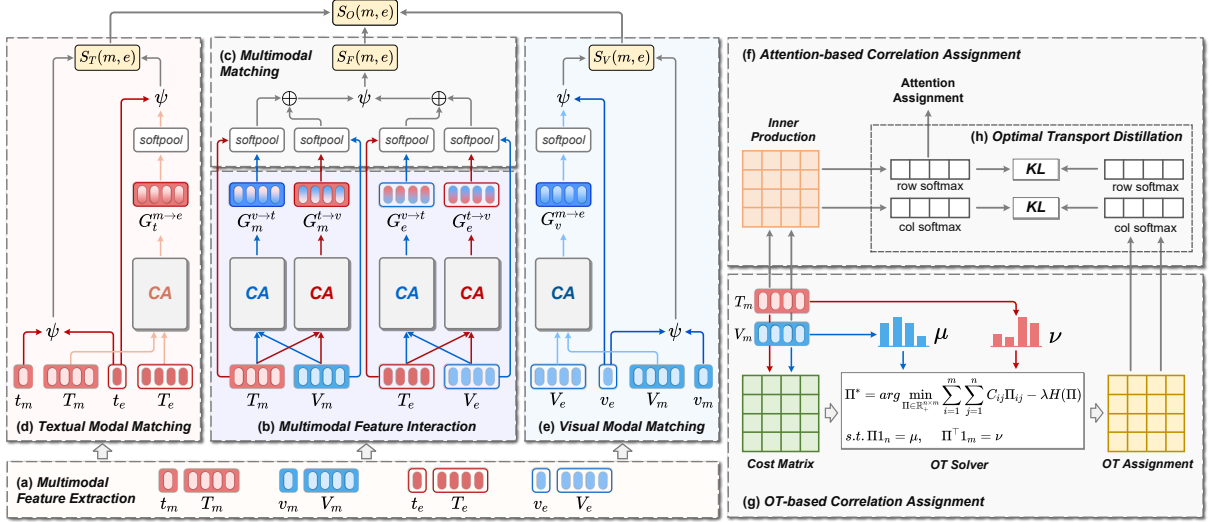


Figure 2: The overall OT-MEL framework. CA denotes the correlation assignment modules with attention or OT.

and $\mathbf{1}_m$ are all-ones vectors of dimension n and m , respectively. In this way, the transport plan remains the multi-element correlations between two groups with appropriate weight assignments.

To accelerate the calculation, Cuturi (2013) proposes to smooth the problem with the entropy regularization term as follows:

$$\begin{aligned} \Pi^* &= \arg \min_{\Pi \in \mathbb{R}_+^{n \times m}} \sum_{i=1}^m \sum_{j=1}^n C_{ij} \Pi_{ij} - \lambda H(\Pi), \\ \text{s.t. } \Pi \mathbf{1}_n &= \mu, \quad \Pi^\top \mathbf{1}_m = \nu, \end{aligned} \quad (3)$$

where $H(\Pi)$ is the entropy of Π , λ is regularization coefficient. Based upon, the OT problem can be efficiently solved by Sinkhorn-Knopp algorithm, detailed in Algorithm 1.

3 Method

In this section, we introduce our OT-MEL model, which consists of three modules: 1) multimodal feature encoder, 2) multimodal feature interaction, and 3) multimodal entity linking. The architecture of OT-MEL is shown in Figure 2.

3.1 Multimodal Feature Extraction

Given the mention m and entity e with their text and image, we are going to generate their initial embeddings. To enrich fundamental correlations between multimodal features, following Luo et al. (2023), we adopt contrastive language-image pre-training architecture (CLIP) (Radford et al., 2021) as the multimodal encoder, which contains a pre-trained BERT (Devlin et al., 2018) for textual embedding and a pre-trained ViT (Dosovitskiy et al.,

2020) for visual embedding. For entities and mentions we use a parameter shared encoder to ensure consistency in the representation space.

Textual Feature Extraction We employ BERT to extract textual features. For mention m , we concatenate the word of mention m_n with its context sentence m_c , and then feed it into BERT as:

$$\mathbf{T}_m = \text{BERT}([\text{CLS}]m_n[\text{SEP}]m_c[\text{SEP}]), \quad (4)$$

where $\mathbf{T}_m \in \mathbb{R}^{L_{mt} \times d}$ is the mention textual embeddings. On the other hand, for entity e , we concatenate the entity name e_n with its attributes³ e_a as $[\text{CLS}]e_n[\text{SEP}]e_a[\text{SEP}]$, and use the same way to obtain the entity textual embeddings $\mathbf{T}_e \in \mathbb{R}^{L_{et} \times d}$.

Visual Feature Extraction We employ the pre-trained ViT to extract visual features. To process the mention image into token sequence, we reshape the image $\mathbf{m}^v \in \mathbb{R}^{H \times W \times C}$ into a sequence of flattened 2D patches as $\mathbf{m}_p^v \in \mathbb{R}^{P \times P \times C}$ with special token [CLS], and feed them into ViT as:

$$\mathbf{V}_m = \text{ViT}([\mathbf{m}_{[\text{CLS}]}, \mathbf{m}_1^v, \dots, \mathbf{m}_{l_v-1}^v]), \quad (5)$$

where $\mathbf{T}_m \in \mathbb{R}^{L_{mv} \times d}$ is the mention visual embeddings, and $l_v = HW/P^2$ is the resulted patch number. Similarly, we reshape the image of entity into patches with the same patch size, and obtain the entity visual embeddings $\mathbf{V}_e \in \mathbb{R}^{L_{ev} \times d}$.

3.2 Multimodal Feature Interaction

As claimed in the introduction, the text and image can have multi-element correlations served as

³Following Luo et al. (2023), we also treat different attributes as texts separated by period.

message passing pivots. Therefore, we assign distinct weights for multimodal feature interaction. Take the mention-side texts and images as examples, we first introduce an intuitive attention-based assignment, and thereupon propose our OT-based correlation assignment method.

Attention-based Correlation Assignment To interact multimodal features with distinct weights, an idea is to employ attention mechanism (Vaswani et al., 2017) between the text and image. Given the mention m with its textual embeddings \mathbf{T}_m and visual embeddings \mathbf{V}_m , we establish the image-to-text correlation assignments as follows:

$$\mathbf{Q} = \mathbf{T}_m \mathbf{W}_q, \mathbf{K} = \mathbf{V}_m \mathbf{W}_k, \mathbf{H} = \mathbf{V}_m \mathbf{W}_v, \quad (6)$$

$$\mathbf{A}^{v \rightarrow t} = \text{softmax}(\mathbf{S}), \mathbf{S} = \mathbf{Q} \mathbf{K}^\top / \sqrt{d} \quad (7)$$

$$\mathbf{G}_m^{v \rightarrow t} = \mathbf{A}^{v \rightarrow t} \mathbf{H} \quad (8)$$

where $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v$ are learnable parameters with $\mathbb{R}^{d \times d}$. Here $\mathbf{A}_{ij}^{v \rightarrow t} \in \mathbb{R}^{L_{mt} \times L_{mv}}$ is the assigned correlation weight from visual patch $\mathbf{V}_{m,j}$ to textual token $\mathbf{T}_{m,i}$. As claimed, the attention-based assignment focuses on local similarity of element-pairs without global consideration.

OT-based Correlation Assignment To globally consider the correlations, we propose OT-based correlation assignment method. Practically, we formulate the image-to-text correlation assignment as an OT problem, where we hope the textual tokens and visual patches has appropriate correlation with minimal total transport cost. Therefore, we project original textual feature \mathbf{T}_m and visual feature \mathbf{V}_m into an assignment feature space, and define the transport cost from patches to tokens with a measurement ϕ as follows:

$$\mathbf{Q} = \mathbf{T}_m \mathbf{W}_q, \mathbf{K} = \mathbf{V}_m \mathbf{W}_k, \mathbf{H} = \mathbf{V}_m \mathbf{W}_h, \quad (9)$$

$$\mathbf{C}_{ij} = \phi(\mathbf{Q}_i, \mathbf{K}_j) \triangleq \frac{1}{2} [1 - \cos(\mathbf{Q}_i, \mathbf{K}_j)], \quad (10)$$

where we achieve ϕ with cosine similarity such that $\phi \in [0, 1]$. In this way, the token-patch pairs with higher similarity would has lower transport cost. Based upon, we treat each element of tokens and patches equally as uniform distributions, and solve the correlation assignment problem by Sinkhorn-Knopp algorithm (Cuturi, 2013):

$$\mathbf{A}^{v \rightarrow t} = \mathbf{\Pi}^* = \text{Sinkhorn-Knopp}(\mathbf{C}, \boldsymbol{\mu}, \boldsymbol{\nu}), \quad (11)$$

$$\mathbf{G}_m^{v \rightarrow t} = \mathbf{A}^{v \rightarrow t} \mathbf{H} \quad (12)$$

where $\boldsymbol{\mu} = 1/L_{mt} \mathbf{1}_{mt}$ and $\boldsymbol{\nu} = 1/L_{mv} \mathbf{1}_{mv}$ are the probabilistic simplexes of textual tokens and visual patches, respectively. In this way, we establish the image-to-text correlations with OT assignment in Eq. (11), which ensures the overall transport cost to be minimal. Actually, this design treats all tokens and patches equally, and ensures all token-patch pairs to be appropriately transported according to the similarity-related cost, thus avoiding the sub-optimal assignment dominated by partial elements. In summary, the above procedure holds for both mentions and entities, and all generated features are termed as $\mathbf{G}_m^{v \rightarrow t}, \mathbf{G}_m^{t \rightarrow v}, \mathbf{G}_e^{t \rightarrow v}$ and $\mathbf{G}_e^{v \rightarrow t}$.

3.3 Multimodal Entity Linking

With the interacted multimodal features, we conduct MEL by matching the mentions and entities. To enhance the interaction between unimodal features of mentions and entities, we also conduct OT-based correlation assignment on them.

Multimodal Feature Matching To fully consider the multimodal information, we integrate the interacted features to match mentions and entities, respectively. Take mention m as an example, we integrate the image-to-text representation with its textual embedding \mathbf{T}_m and the transported embedding $\mathbf{G}_m^{v \rightarrow t}$ from its image. Generally, we can employ any pooling method to aggregate these embeddings as $\mathcal{G} = \{\mathbf{T}_m, \mathbf{G}_m^{v \rightarrow t}\}$. Here we introduce a simple yet effective method, softpool (Stergiou et al., 2021), which is a soft version of max-pooling:

$$\bar{\mathbf{g}}_m^t = \text{softpool}(\mathcal{G}) \triangleq \sum_{\mathbf{g}_i \in \mathcal{G}} \mathbf{w}_i \odot \mathbf{g}_i, \quad (13)$$

$$\mathbf{w}_i = \frac{\exp(\mathbf{g}_i)}{\sum_{\mathbf{g}_j \in \mathcal{G}} \exp(\mathbf{g}_j)} \quad (14)$$

where $\bar{\mathbf{g}}_m^t \in \mathbb{R}^d$ is the generated textual representation, \odot is Hadamard production. In this way, it emphasizes the most activated features in the generated multimodal representation.

Based upon, we obtain the text-to-image representations of mention m , and those representations of entity e as follows:

$$\bar{\mathbf{g}}_m^v = \text{softpool}(\{\mathbf{V}_m, \mathbf{G}_m^{t \rightarrow v}\}) \quad (15)$$

$$\bar{\mathbf{g}}_e^t = \text{softpool}(\{\mathbf{T}_e, \mathbf{G}_e^{v \rightarrow t}\}) \quad (16)$$

$$\bar{\mathbf{g}}_e^v = \text{softpool}(\{\mathbf{V}_e, \mathbf{G}_e^{t \rightarrow v}\}) \quad (17)$$

Finally, we measure the matching scores between mentions and entities, considering both of their

textual and visual representations:

$$S_F(m, e) = \psi(\bar{\mathbf{g}}_m^t \oplus \bar{\mathbf{g}}_m^v, \bar{\mathbf{g}}_e^t \oplus \bar{\mathbf{g}}_e^v) \quad (18)$$

where ψ is the similarity measurement, and we achieve it by vector inner production. $S_F(m, e)$ is the fused multimodal feature matching score.

Unimodal Feature Matching Considering the unimodal correlations between mentions and entities, we design unimodal feature matching module with OT-based correlation assignment. Specifically, take the text-modal as an example, we capture the correlations between mention textual feature \mathbf{T}_m to entity \mathbf{T}_e . We also derive the cost matrix with cosine similarity between them, and obtain the textual modal correlated mention representation:

$$\mathbf{Q} = \mathbf{T}_e \mathbf{W}'_q, \mathbf{K} = \mathbf{T}_m \mathbf{W}'_k, \mathbf{H} = \mathbf{T}_m \mathbf{W}'_v, \quad (19)$$

$$C_{ij} = \phi(\mathbf{Q}_i, \mathbf{K}_j) \triangleq \frac{1}{2} [1 - \cos(\mathbf{Q}_i, \mathbf{K}_j)], \quad (20)$$

$$\mathbf{A}^{m \rightarrow e} = \mathbf{\Pi}^* = \text{Sinkhorn-Knopp}(\mathbf{C}, \boldsymbol{\mu}, \boldsymbol{\nu}), \quad (21)$$

$$\mathbf{G}_t^{m \rightarrow e} = \mathbf{A}^{m \rightarrow e} \mathbf{H} \quad (22)$$

where $\boldsymbol{\mu} = 1/L_{mt} \mathbf{1}_{mt}$ and $\boldsymbol{\nu} = 1/L_{et} \mathbf{1}_{et}$ are the probabilistic simplexes. Based upon, we aggregate the textual representations for the mention, and achieve the unimodal feature matching between mention and entity by:

$$\bar{\mathbf{g}}_m = \text{softpool}(\{\mathbf{G}_t^{m \rightarrow e}\}) \quad (23)$$

$$S_T(m, e) = \frac{1}{2} [\psi(\bar{\mathbf{g}}_m, \mathbf{t}_e) + \psi(\mathbf{t}_m, \mathbf{t}_e)] \quad (24)$$

where we also use inner production to achieve ψ . Here, \mathbf{t}_m and \mathbf{t}_e are the output embedding of token [CLS], remaining the original summarized feature of texts. This procedure also holds for visual modal feature matching with matching score $S_V(m, e)$.

Overall Matching Score To consider different aspects in feature matching, we collect all matching scores and obtain the overall matching score as:

$$S_O(m, e) = \frac{1}{3} \sum_{X \in \{F, T, V\}} S_X(m, e) \quad (25)$$

Thereby, we design contrastive training objective on the overall matching score as follows:

$$\mathcal{L}_O = -\log \frac{\exp(S_O(m, e))}{\sum_{e' \in \mathcal{E}} \exp(S_O(m, e'))}, \quad (26)$$

where \mathcal{E} is all candidate entities in the KG. For training, we adopt all other entity e' except the true

entity e in the batch, which serve as the in-batch negative samples. To enhance the consistency in different matching aspects, following Luo et al. (2023), we also apply contrastive training objective on different matching scores:

$$\mathcal{L}_X = -\log \frac{\exp(S_X(m, e))}{\sum_{e' \in \mathcal{E}} \exp(S_X(m, e'))}, \quad (27)$$

$$\mathcal{L} = \mathcal{L}_O + \sum_{X \in \{F, T, V\}} \mathcal{L}_X$$

where $X \in \{F, T, V\}$ is the original multimodal and unimodal matching scores. \mathcal{L} is the overall loss of our model with the full OT-based version.

3.4 Optimal Transport Distillation

To improve the efficiency of OT-based correlation assignments, we further propose an efficient version with KD, which transfers knowledge from OT-based correlation assignment to attention-based correlation assignment. Generally, we can use mean squared error (MSE) loss to assimilate the assignment $\mathbf{A}^{v \rightarrow t}$ in Eq. (7) and Eq. (11). To better consider both the source and target aspects in OT assignment, we design Kullback-Leibler (KL) divergence to fully transfer knowledge as follows:

$$\text{KD}(\mathbf{A}_{\text{OT}}^{v \rightarrow t}, \mathbf{A}_{\text{ATT}}^{v \rightarrow t})$$

$$\triangleq \frac{1}{2} [\sum_j \text{KL}(\text{softmax}(\mathbf{\Pi}_{i*}^*), \text{softmax}(\mathbf{S}_{i*}))$$

$$+ \sum_j \text{KL}(\text{softmax}(\mathbf{\Pi}_{*j}^*), \text{softmax}(\mathbf{S}_{*j}))] \quad (28)$$

Here, we conduct knowledge distillation on all OT-based correlation assignments, including $\mathbf{A}^{m, v \rightarrow t}$, $\mathbf{A}^{m, t \rightarrow v}$, $\mathbf{A}^{e, v \rightarrow t}$, $\mathbf{A}^{e, t \rightarrow v}$ for multimodal feature interaction, and $\mathbf{A}^{m \rightarrow e}$, $\mathbf{A}^{e \rightarrow m}$ for unimodal feature matching. All these knowledge distillation constraints constitute the overall KD loss \mathcal{L}_{KD} . Alternatively, the overall training loss is:

$$\mathcal{J} = \mathcal{L} + \mathcal{L}_{\text{KD}} \quad (29)$$

Here, \mathcal{J} is the overall loss with knowledge distillation, which obtains an efficient KD-version model.

4 Experiments

In this section, we carried out comprehensive experiments on three public MEL benchmark datasets to evaluate the effectiveness of OT-MEL model.

4.1 Datasets

In the experiments, we selected three public MEL datasets RichpediaMEL with 17K samples,

Model	RichpediaMEL				WikiMEL				WikiDiverse			
	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR
BLINK	58.47	81.51	88.09	71.39	74.66	86.63	90.57	81.72	57.14	78.04	85.32	69.15
BERT	59.55	81.12	87.16	71.67	74.82	86.79	90.47	81.78	55.77	75.73	83.11	67.38
RoBERTa	61.34	81.56	87.15	72.80	73.75	85.85	89.80	80.86	59.46	78.54	85.08	70.52
CLIP	67.78	85.22	90.04	77.57	83.23	92.10	94.51	88.23	61.21	79.63	85.18	71.69
VILT	45.85	62.96	69.80	56.63	72.64	84.51	87.86	79.46	34.39	51.07	57.83	45.22
ALBEF	65.17	82.84	88.28	75.29	78.64	88.93	91.75	84.56	60.59	75.59	81.30	69.93
METER	63.96	82.24	87.08	74.15	72.46	84.41	88.17	79.49	53.14	70.93	77.59	63.71
DZMNED	68.16	82.94	87.33	76.63	78.82	90.02	92.62	84.97	56.90	75.34	81.41	67.59
JMEL	48.82	66.77	73.99	60.06	64.65	79.99	84.34	73.39	37.38	54.23	61.00	48.19
VELML	67.71	84.57	89.17	77.19	76.62	88.75	91.96	83.42	54.56	74.43	81.15	66.13
GHMFC	72.92	86.85	90.60	80.76	76.55	88.40	92.01	83.36	60.27	79.40	84.74	70.99
MIMIC	81.02	91.77	94.38	86.95	87.98	95.07	96.37	91.82	63.51	81.04	86.43	73.44
ATT	78.33	91.16	94.33	85.28	88.20	94.87	96.52	91.93	63.38	81.95	<u>87.25</u>	73.82
OT-MEL (<i>OT</i>)	83.30	92.39	94.83	88.27	88.97	95.63	96.96	92.59	66.07	82.82	87.39	75.43
OT-MEL (<i>KD</i>)	<u>82.57</u>	<u>92.06</u>	<u>94.44</u>	<u>87.78</u>	<u>88.37</u>	<u>95.41</u>	<u>96.90</u>	<u>92.26</u>	<u>64.82</u>	<u>82.29</u>	<u>87.25</u>	<u>74.65</u>

Table 1: Experimental results (%) on three MEL benchmark datasets, including RichpediaMEL, WikiMEL and WikiDiverse. The **bold** scores are the best results, and the underline scores are the second results.

WikiMEL (Wang et al., 2022a) with 25K samples and WikiDiverse (Wang et al., 2022b) with 15K samples to verify the effectiveness of our proposed method. We follow the training, validation, and testing set splitting approach as Luo et al. (2023). Following Wang et al. (2022a) we use the subset of Wikidata as KG for each dataset. The detail of three datasets information are reported in Appendix C.

4.2 Baselines

We compare our method with various competitive baselines in three groups: 1) the text-based methods, which utilize textual information only to achieve EL, including **BLINK** (Wu et al., 2019), **BERT** (Devlin et al., 2018), and **RoBERTa** (Liu et al., 2019). 2) the vision-and-language pre-training (VLP) methods including **CLIP** (Radford et al., 2021), **VILT** (Kim et al., 2021), **ALBEF** (Li et al., 2021), and **METER** (Dou et al., 2022). 3) the MEL enhanced methods⁴, which explore sophisticated interaction networks to capture diverse correlations, including **DZMNED** (Moon et al., 2018), **JMEL** (Adjali et al., 2020), **VELML** (Zheng et al., 2022), **GHMFC** (Wang et al., 2022a), **MIMIC** (Luo et al., 2023). For details, we provide extensive descriptions of baselines in Appendix B.

4.3 Implementation Details

Our model weights are initialized with pre-trained CLIP-Vit-Base-Patch32. We set the batch size to

⁴For fairness, we don’t compare with DRIN (Xing et al., 2023), since we fail to run the model due to the numerous scene graphs of entities.

256, the learning rate to 2×10^{-5} , and the hidden dimension d to 96. For the OT parameter λ , we set it to 0.6. All methods are evaluated on the validation set and the checkpoint with the highest MRR is selected to evaluate on the test set.

4.4 Main Results

We conduct comparative experiments on three public datasets to verify the effectiveness of each model on MEL tasks, Table 1 shows the H@k and MRR of different models in the three datasets, and we will provide the formula for metrics in the Appendix D. Based on the experimental results, we further have the following observations and analysis: (1) *Multimodal information is important for EL problems.* The first section of the table presents text-based EL models, which, despite demonstrating commendable performance, significantly lag behind the multimodal models showcased in the second section of the table. (2) *Our model is effective on all three datasets.* Compared with the state-of-the-art model, OT-MEL improves H@1 by 2.28%, 2.58% and 0.99% on the three datasets, respectively. (3) *OT can indeed guide the model for better modality fusion.* The results of OT-MEL(KD) based on distillation on 3 datasets also exceed the current state-of-the-art models.

4.5 Variant Analysis

To inspect the effectiveness of the model components, we show results of experiments for model variants in Table 2, where we find: (1) *OT is more effective than attention mechanisms in both uni-*

No.	Variant	H@1	MRR	Δ Avg
-	OT-MEL (OT)	83.30	88.27	-
F1	repl. ATT	80.60	86.70	\downarrow 2.14
F2	w/o FusM	78.94	85.42	\downarrow 3.61
M1	repl. ATT	81.72	87.61	\downarrow 1.12
M2	w/o UniM	73.95	81.85	\downarrow 7.89
A1	ATT (w/o OT)	78.33	85.28	\downarrow 3.98
O1	SoftPool \rightarrow Mean	80.80	86.55	\downarrow 2.11
O2	SoftPool \rightarrow Max	82.28	87.25	\downarrow 1.02
-	OT-MEL (KL)	82.57	87.78	-
K1	KD \rightarrow MSE	81.95	87.34	\downarrow 0.53
K2	KD \rightarrow UniKL	81.84	87.33	\downarrow 0.59
K3	ATT (w/o KD)	78.33	85.28	\downarrow 3.37

Table 2: Results (%) of variants on the Richpediamel data. Δ Avg is the average decrease in MRR and H@1. The details and full results are reported in Appendix E.

modal and multimodal feature matching. In models F1 and M1, where we replaced OT with attention, there was a noticeable decrease in performance for both. (2) *Multimodal feature fusion can better utilize the complementary information between modalities to aid in linking.* In model F2, after removing the multimodal feature fusion module, the average results decreased by 3.61%. (3) *Unimodal feature matching can more effectively utilize the fine-grained correlations between images and texts to enhance EL.* In model M2, when the unimodal feature matching module is removed, the performance decreases by 7.89%. (4) *Compared to mean pooling and max pooling, SoftPool can aggregate sequential data more effectively.* After replacing SoftPool with mean pooling and max pooling, the performance decreased by 2.11% and 1.02%, respectively. (5) *OT-MEL(KL) can enhance the inference speed of the model while only marginally compromising model performance.* OT-MEL(KL), compared to OT-MEL(OT), improves the inference speed by 337.20s on the validation set of the Richpediamel dataset, while the performance only decreases by 0.62%. Table 4 displays the average inference speeds on the validation sets of different datasets.

4.6 Analysis on MEL with Multi-objects

To validate that our model can better utilize the multi-element correlations between different modalities to aid in EL, we used YOLOv5 for object detection on images in the Richpediamel test dataset. We categorized images into multi-object sets for object counts of two or more, and single-




Figure 3: Impact of entropy $H(\Pi)$ on Richpediamel.

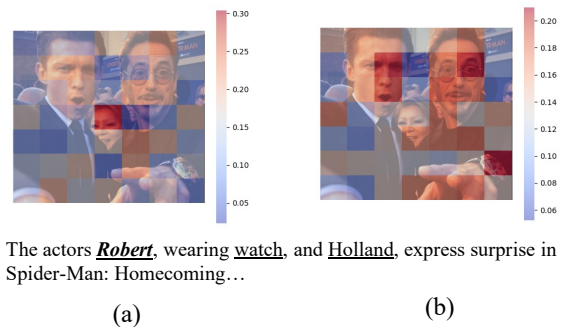


Figure 4: Case study of multi-element correlation assignments: (a) attention matrix, (b) OT matrix.

object sets for counts of one or fewer. Notably, we excluded data where EL was straightforward due to identical mentions and entity names with no ambiguous counterparts in the KG, ensuring test rigor. Table 3 shows the results of our experiment. Observations indicate that *our model can indeed achieve better linking by utilizing the multi-element correlations between multimodalities.* Compared to the attention mechanism, our model shows a 6.19% improvement in H@1 on single-object sets and a 10.40% improvement in H@1 on multi-object sets.

4.7 Case Study of Optimal Transport

The entropy controls the information dispersion in the transport plan of OT, where the higher entropy tends to have a more disperse concentration. Figure 3 shows H@1 and MRR results. The observation reveals that *an appropriate entropy allows the model to better focus on the information it should.* We can see that when the entropy is low, the results are poor due to an overly concentrated focus. Besides, when the entropy is high, the results tend to be worse due to an overly concentrated focus.

5 Related Works

Text-based Entity Linking Traditional EL mainly focuses on the textual information of mentions and entities. These methods (Cao et al., 2017; Wu et al., 2019) first employ textual encoders to

Model	Single-object Image				Multi-object Image				Overall			
	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR
MIMIC	75.10	87.82	91.46	82.34	79.41	92.78	95.97	86.50	75.85	88.69	92.28	83.07
ATT	76.09	90.51	93.71	83.81	75.16	91.08	95.12	83.82	71.77	88.43	92.58	80.81
OT-MEL (<i>OT</i>)	82.28	92.16	94.87	87.72	85.56	94.69	96.18	90.31	78.82	90.17	93.21	85.02

Table 3: Multi-element correlation analysis on Richpediamel according to images with different object numbers.

Model	Richpediamel	Wikimel	Wikidiverse
MIMIC	1,033.47s	1,000.12s	685.69s
ATT	1,093.03s	1,875.24s	1,543.87s
OT-MEL (<i>KD</i>)	1,842.96s	2,012.95s	1,595.17s
OT-MEL (<i>OT</i>)	2,175.66s	3,028.11s	1,841.67s

Table 4: Efficiency comparison (s) in prediction.

obtain contextual representations of mention and entity, and then calculate similarities between them as final probabilities. To enhance consistency in EL, they enrich local correlations (Peters et al., 2019; Liu et al., 2023) with prior knowledge and collective mechanism, and global correlations (Le and Titov, 2018; Cao et al., 2018; Fang et al., 2019) with topical coherence among mentions. However, these methods are designed to deal with text only, which can hardly deal with the crucial multimodal correlations within images in MEL. Other studies have focused on Generative Entity Linking (GEL) (De Cao et al., 2020; Shi et al., 2023; De Cao et al., 2022; Huang et al., 2022; Yuan et al., 2022). This is a promising research problem in the era of Large Language Models (LLMs), but they can suffer from the lexical ambiguity problem (Fortuny and Payrató, 2024), which we will discuss in details in the Appendix G.

Multimodal Entity Linking Recent years have seen numerous social media and news posted with texts and images, thereby MEL is proposed to identify interested entities from multimodal information, which can benefit to knowledge acquisition (Sheng et al., 2020; Li et al., 2023). As a pioneering study, Moon et al. (2018) firstly introduces images to benefit EL, which blends visual, word and character features through an attention mechanism. Afterward, there exist several studies (Adjali et al., 2020; Zheng et al., 2022; Wang et al., 2022a; Luo et al., 2023) exploring gate or attention mechanism to capture complicated multimodal correlations for mentions and entities. Recently, Xing et al. (2023) introduce external object-level scene graphs of images to enrich visual correlations, but may

involve elaborate errors and still construct multimodal correlations with attentive graph neural networks. Luo et al. (2023) present a sophisticated gated framework for fine-grained inter-modal interactions, achieving the most advanced results. Although existing studies attempt multiple correlations in different aspects, they heavily rely on the automatically learned attention-based mechanism, which can hardly ensure appropriate correlation assignments and impede further improvements.

Optimal Transport OT (Kantorovich, 1942) is a well-known problem, which can trace back to Monge (1781). The key idea is to derive an optimal transport plan to transfer one distributions to another (Arjovsky et al., 2017; Chen et al., 2020; Bhardwaj et al., 2021; Yang et al., 2023). There are extensive studies utilizing the transport plan of OT to solve assignment problems (Chang et al., 2022; Cao et al., 2022). Liu et al. (2020b) employ the global OT plan to address the many-to-one issue in semantic correspondence in computer vision. Lee et al. (2019); Chen et al. (2020) employ the wasserstein distance to align the representation of text and image. However, few existing studies conduct comprehensive investigation of the OT plan in the context of multi-element correlations within both unimodal and multimodal settings.

6 Conclusion

This paper introduces a MEL model with OT-guided correlation assignment, termed OT-MEL. Previous studies rely heavily on the automatically learned attention weights, which may over-concentrate on partial correlations. To this end, OT-MEL proposes capturing these multi-element correlations across modalities with OT. To improve the efficiency, we further transfer knowledge from OT to attention mechanism. Experiments on 3 datasets illustrate that our model outperforms current state-of-the-art models. In the future, we will explore other efficient method to deploy OT in MEL.

Limitations

In this paper, we propose a MEL model with OT correlation assignment, namely OT-MEL. First, OT-MEL simplistically assumes uniform distributions for both source and target, which can be further explored in real-world scenarios. Besides, as OT-MEL is tailored for visual and textual modality in MEL, it doesn't consider OT in other modalities such as speech and video. We will explore the generality of the proposed OT correlation assignment in these complex modalities for wide research communities.

Acknowledgement

We would like to thank the anonymous reviewers for their comments. This work was supported by the National Key Research and Development Program of China (Grant No.2021YFB3100600), the Youth Innovation Promotion Association of CAS (No.2021153), and the Postdoctoral Fellowship Program of CPSF (No.GZC20232968).

References

- Omar Adjali, Romaric Besançon, Olivier Ferret, Hervé Le Borgne, and Brigitte Grau. 2020. Multimodal entity linking for tweets. In *European Conference on Information Retrieval*, pages 463–478. Springer.
- Yuki Arase, Han Bao, and Sho Yokoi. 2023. Unbalanced optimal transport for unbalanced word alignment. *arXiv preprint arXiv:2306.04116*.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR.
- Rishabh Bhardwaj, Tushar Vaidya, and Soujanya Poria. 2021. Knot: Knowledge distillation using optimal transport for solving nlp tasks. *arXiv preprint arXiv:2110.02432*.
- Yixin Cao, Lei Hou, Juanzi Li, and Zhiyuan Liu. 2018. Neural collective entity linking. *arXiv preprint arXiv:1811.08603*.
- Yixin Cao, Lifu Huang, Heng Ji, Xu Chen, and Juanzi Li. 2017. Bridge text and knowledge by learning multi-prototype entity mention embedding.
- Zongsheng Cao, Qianqian Xu, Zhiyong Yang, Yuan He, Xiaochun Cao, and Qingming Huang. 2022. Otkge: Multi-modal knowledge graph embeddings via optimal transport. *Advances in Neural Information Processing Systems*, 35:39090–39102.
- Wanxing Chang, Ye Shi, Hoang Tuan, and Jingya Wang. 2022. Unified optimal transport framework for universal domain adaptation. *Advances in Neural Information Processing Systems*, 35:29512–29524.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer.
- Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2020. Autoregressive entity retrieval. *arXiv preprint arXiv:2010.00904*.
- Nicola De Cao, Ledell Wu, Kashyap Papat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022. Multilingual autoregressive entity linking. *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. 2022. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18166–18176.
- Zheng Fang, Yanan Cao, Qian Li, Dongjie Zhang, Zhenyu Zhang, and Yanbing Liu. 2019. Joint entity linking with deep reinforcement learning. In *The world wide web conference*, pages 438–447.
- Jordi Fortuny and Lluís Payrató. 2024. Ambiguity in linguistics 1. *Studia Linguistica*, 78(1):1–7.
- Guanhuan Huang, Xiaojun Quan, and Qifan Wang. 2022. Autoregressive entity generation for end-to-end task-oriented dialog. *arXiv preprint arXiv:2209.08708*.
- Jiyue Jiang, Sheng Wang, Qintong Li, Lingpeng Kong, and Chuan Wu. 2023. A cognitive stimulation dialogue system with multi-source knowledge fusion for elders with cognitive impairment. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*,

- pages 10628–10640, Toronto, Canada. Association for Computational Linguistics.
- Leonid V Kantorovich. 1942. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pages 199–201.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR.
- Phong Le and Ivan Titov. 2018. Improving entity linking by modeling latent relations between mentions. *arXiv preprint arXiv:1804.10637*.
- John Lee, Max Dabagia, Eva Dyer, and Christopher Rozell. 2019. Hierarchical optimal transport for multimodal distribution alignment. *Advances in neural information processing systems*, 32.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705.
- Qian Li, Shu Guo, Yangyifei Luo, Cheng Ji, Lihong Wang, Jiawei Sheng, and Jianxin Li. 2023. Attribute-consistent knowledge graph representation learning for multi-modal entity alignment. In *Proceedings of the ACM Web Conference 2023*, pages 2499–2508.
- Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. 2020a. Semantic correspondence as an optimal transport problem. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 4462–4471. Computer Vision Foundation / IEEE.
- Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. 2020b. Semantic correspondence as an optimal transport problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4463–4472.
- Yi Liu, Yuan Tian, Jianxun Lian, Xinlong Wang, Yanan Cao, Fang Fang, Wen Zhang, Haizhen Huang, Denvy Deng, and Qi Zhang. 2023. Towards better entity linking with multi-view enhanced distillation. *arXiv preprint arXiv:2305.17371*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2018. Entity-duet neural ranking: Understanding the role of knowledge graph semantics in neural information retrieval. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2395–2405, Melbourne, Australia. Association for Computational Linguistics.
- Pengfei Luo, Tong Xu, Shiwei Wu, Chen Zhu, Linli Xu, and Enhong Chen. 2023. Multi-grained multimodal interaction network for entity linking. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1583–1594.
- Gaspard Monge. 1781. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704.
- 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.
- Matthew E Peters, Mark Neumann, Robert L Logan IV, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. 2019. Knowledge enhanced contextual word representations. *arXiv preprint arXiv:1909.04164*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Jiawei Sheng, Shu Guo, Zhenyu Chen, Juwei Yue, Lihong Wang, Tingwen Liu, and Hongbo Xu. 2020. Adaptive attentional network for few-shot knowledge graph completion. *arXiv preprint arXiv:2010.09638*.
- Senbao Shi, Zhenran Xu, Baotian Hu, and Min Zhang. 2023. Generative multimodal entity linking. *arXiv preprint arXiv:2306.12725*.
- Alexandros Stergiou, Ronald Poppe, and Grigorios Kalliatakis. 2021. Refining activation downsampling with softpool. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10357–10366.
- 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, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural*

- Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Cédric Villani et al. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Meng Wang, Haofen Wang, Guilin Qi, and Qiushuo Zheng. 2020. Richpedia: a large-scale, comprehensive multi-modal knowledge graph. *Big Data Research*, 22:100159.
- 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*, pages 938–948.
- Xuwu Wang, Junfeng Tian, Min Gui, Zhixu Li, Rui Wang, Ming Yan, Lihan Chen, and Yanghua Xiao. 2022b. Wikidiverse: a multimodal entity linking dataset with diversified contextual topics and entity types. *arXiv preprint arXiv:2204.06347*.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2019. Scalable zero-shot entity linking with dense entity retrieval. *arXiv preprint arXiv:1911.03814*.
- Shangyu Xing, Fei Zhao, Zhen Wu, Chunhui Li, Jianbing Zhang, and Xinyu Dai. 2023. Drin: Dynamic relation interactive network for multimodal entity linking. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3599–3608.
- Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. Improving question answering over incomplete KBs with knowledge-aware reader. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4258–4264, Florence, Italy. Association for Computational Linguistics.
- Wei Yang, Jie Yang, and Yuan Liu. 2023. Multimodal optimal transport knowledge distillation for cross-domain recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 2959–2968.
- Hongyi Yuan, Zheng Yuan, and Sheng Yu. 2022. Generative biomedical entity linking via knowledge base-guided pre-training and synonyms-aware fine-tuning. *arXiv preprint arXiv:2204.05164*.
- Qiushuo Zheng, Hao Wen, Meng Wang, and Guilin Qi. 2022. Visual entity linking via multi-modal learning. *Data Intelligence*, 4(1):1–19.

A Optimal Transport




Figure 5: The differences between Monge and Kantorovich formulations in optimal transport, with a multimodal problem example.

Optimal Transport (OT) is a problem which traces back to the work of Gaspard Monge (Monge, 1781). It is a constrained optimization problem that seeks an efficient solution of transporting one distribution of mass to another. Let μ and ν be two probability measures defined on the same space \mathbb{R}^n , representing the source and target distributions of resources, respectively. The Monge problem aims to find a mapping $\Pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that transforms a distribution μ into a distribution ν while minimizing the total transfer cost. This cost is usually measured by a cost function $c(x, \Pi(x))$, where x is the original location of the resource and $\Pi(x)$ is the new location to which it has been transferred. It can be formulated as the following optimization problem:

$$\Pi^* = \arg \min_{\Pi} \int_{\mathbb{R}^n} c(x, \Pi(x)) d\mu(x), \quad (30)$$

where Π is the transport map, $c(x, \Pi(x))$ is a cost function, measure from the x to the $\Pi(x)$ transfer cost, and the integral is the total cost over all x .

However, there is a very strong assumption in the Monge problem which is the inseparability of masses, that is, the Monge problem requires an explicit mapping Π to transfer directly from one location to another, and when the two distributions have different total masses or there is a many-to-many mapping between the source distribution and the target distribution, the Monge problem cannot find a feasible solution. For example, in multimodal problem, the total mass in the image distribution and the total mass in the text distribution may be different, and an image patch may correspond to multiple text tokens. Kantorovich (Kantorovich, 1942) proposed an extension and generalization of the Monge problem to address the aforementioned limitations in the Monge problem. Figure 5

illustrates the difference between the Monge Formulation and the Kantorovich Formulation.

Algorithm 1: Sinkhorn-Knopp Algorithm

Require: Cost matrix $\mathbf{C} \in \mathbb{R}^{n \times m}$, marginals $\mu \in \mathbb{R}^n$, $\nu \in \mathbb{R}^m$, regularization parameter $\lambda > 0$

Ensure: Approximate transport matrix \mathbf{P}

Initialize $\mathbf{u} \leftarrow \mathbf{1}_n$, $\mathbf{v} \leftarrow \mathbf{1}_m$

Compute $K \leftarrow \exp(-\lambda \mathbf{C})$

while not converged **do**

 Update $\mathbf{u} \leftarrow \mu \odot (K\mathbf{v})$

 Update $\mathbf{v} \leftarrow \nu \odot (K^T\mathbf{u})$

end while

Compute $\mathbf{P} \leftarrow \text{diag}(\mathbf{u})K\text{diag}(\mathbf{v})$

return \mathbf{P}

Instead of finding a specific transport map, the Kantorovich problem looks for a transport plan, which is a more flexible notion. In this problem, we consider two given probability distributions μ and ν , defined on two spaces X and Y , respectively. The Kantorovich problem aims to find a joint probability distribution (transport plan) Π over $X \times Y$ whose marginal distributions are μ and ν , respectively, while minimizing the overall transport cost. Mathematically, this problem can be expressed as follows:

$$\begin{aligned} \Pi^* &= \arg \min_{\Pi} \int_{X \times Y} c(x, y) \Pi(x, y) dx dy, \\ \text{s.t. } &\int_Y \Pi(x, y) dy = \mu, \int_X \Pi(x, y) dx = \nu, \end{aligned} \quad (31)$$

where Π is a joint probability measure on $X \times Y$, representing the transition plan, and $c(x, y)$ is a cost function for moving from $x \in X$ to $y \in Y$.

Within the scope of our problem, similar to the previous optimal transport (OT) settings (Arase et al., 2023; Chang et al., 2022), we simply assume the tokens of the text and the patches of the image as two different uniform distributions, which means that each token or patch has the same amount of information. Each token or patch is a sample from a discrete probability distribution. So we need to use the Kantorovich problem in discrete form to compute the optimal transport plan between text and image. Specifically, let $X = \{x_1, x_2, \dots, x_n\}$ be the token set of text and $Y = \{y_1, y_2, \dots, y_m\}$ be the patch set of image. The discrete form of Kantorovich problem can be formulated as finding a "transport plan" Π , which is a matrix $n \times m$,

where each element Π_{ij} represents the amount of information to be transported from text token x_i to image patch y_j . This plan needs to minimize the total transportation cost while satisfying constraints on the supply and demand of information. Mathematically, this problem can be formulated as the following linear programming problem:

$$\begin{aligned} \Pi^* &= \arg \min_{\Pi \in \mathbb{R}_+^{n \times m}} \sum_{i=1}^m \sum_{j=1}^n c_{ij} \Pi_{ij}, \\ \text{s.t. } \Pi \mathbf{1}_n &= \boldsymbol{\mu}, \quad \Pi^\top \mathbf{1}_m = \boldsymbol{\nu}, \end{aligned} \quad (32)$$

where Π^* is the optimal transport plan, Π_{ij} is the amount of information from x_i to y_j in the optimal transport plan, $\mathbb{R}_+^{n \times m}$ denotes the space of $n \times m$ real matrices with all non-negative elements, c_{ij} denotes the unit transport cost from x_i to y_j . $\mathbf{1}_n$ is an all-ones column vector of dimension n , and $\mathbf{1}_m$ is an all-ones column vector of dimension m . Since the computation of the Kantorovich problem involves the resolution of a linear program, its cost is prohibitive when dealing with large-scale data. Therefore, [Cuturi \(2013\)](#) proposes to smooth the Kantorovich problem with the entropy regularization term:

$$\begin{aligned} \Pi^* &= \arg \min_{\Pi \in \mathbb{R}_+^{n \times m}} \sum_{i=1}^m \sum_{j=1}^n c_{ij} \Pi_{ij} - \frac{1}{\lambda} H(\Pi), \\ \text{s.t. } \Pi \mathbf{1}_n &= \boldsymbol{\mu}, \quad \Pi^\top \mathbf{1}_m = \boldsymbol{\nu}, \end{aligned} \quad (33)$$

where $H(\Pi)$ is the entropy of Π , $\lambda > 0$ is regularization coefficient, then solve the optimal transport problem quickly by Sinkhorn-Knopp algorithm. Algorithm 1 shows the flow of Sinkhorn-Knopp algorithm, where \oslash is element-wise division.

B Description of Baselines

We compared our method with various competitive baselines. 1) The text-based methods include:

- **BLINK** ([Wu et al., 2019](#)) a two-stage zero-shot EL method, utilizes BERT as its backbone. It first retrieves entities using a bi-encoder, then re-ranks them with a cross-encoder.
- **BERT** ([Devlin et al., 2018](#)) consists of a stack of Transformer encoders and is pre-trained on a large amount of corpus.
- **RoBERTa** ([Liu et al., 2019](#)) further improves BERT by removing next sentence prediction objective and using a dynamic mask strategy.

2) The Vision-and-Language Pre-training (VLP) methods include:

- **CLIP** ([Radford et al., 2021](#)) uses two Transformer-based encoders for visual and textual representation, pre-trained on massive noisy web data using contrastive loss.
- **ViLT** ([Kim et al., 2021](#)) adopts shallow embeddings for text and visuals, focusing on deep modality interaction through a series of Transformer layers.
- **ALBEF** ([Li et al., 2021](#)) aligns visual and textual features using imagetext contrastive loss first, fusing them via a multimodal Transformer encoder, and employs momentum distillation for enhanced learning from noisy data.
- **METER** ([Dou et al., 2022](#)) employs co-attention for semantic interplay between modalities, featuring layers with self-attention, cross-attention modules, and a feedforward network.

3) The MEL methods include:

- **DZMNED** ([Moon et al., 2018](#)) is a trailblazing MEL method, blending visual, word, and character features through an attention mechanism.
- **JMEL** ([Adjali et al., 2020](#)) utilizes unigram and bigram embeddings for text analysis, merging these features via concatenation and a dense layer. For fair comparison, its textual encoder is substituted with a pre-trained BERT model.
- **VELML** ([Zheng et al., 2022](#)) employs the VGG-16 network for object-level visual features and replaces its GRU textual encoder with pre-trained BERT, fusing both modalities using an enhanced attention mechanism.
- **GHMFC** ([Wang et al., 2022a](#)) introduces a hierarchical cross-attention approach for detailed correlation analysis between text and visual features, utilizing contrastive learning for optimization.
- **MIMIC** ([Luo et al., 2023](#)) presents a framework for fine-grained inter-modal interactions, optimized using contrastive learning.

C Statistics of Datasets

WikiMEL ([Wang et al., 2022a](#)) is collected from Wikipedia entities pages and contains more than 22k multimodal sentences. **RichpediaMEL** ([Wang](#)




Figure 6: Experiments on samples with different token counts in the Richpediamel dataset.

et al., 2022a) is obtained from a MMKG Richpedia (Wang et al., 2020). The authors of RichpediaMEL first extracted entities from Richpedia and then obtain multimodal information from Wikidata (Wang et al., 2022a). The main entity types of WikiMEL and RichpedaiMEL are person. WikiDiverse (Wang et al., 2022b) is constructed from Wikinews and covers various topics including sports, technology, economy and so on. Table 5 shows the detail of three datasets. Note that we find there also exist some different entities with the same name, and we report the number of them in the Same Name row in Table 5. Figure 7 displays statistics on the number of text tokens and the count of objects in images across three datasets.

	Richpediamel	Wikimel	Wikidiverse
Train Mentions	12,463	18,092	11,351
Valid Mentions	1,780	2,585	1,664
Test Mentions	3,562	5,169	2,078
Mention Imgs	15,853	22,136	6,697
Same Name	15,514	616	14,827
KG Entities	160,933	109,976	132,460
Entity with Imgs	86,769	67,195	67,309

Table 5: Statistics of three datasets.

D Metrics

We first calculate the similarity scores between a mention and all entities of the KG, then the similarity scores are sorted in descending order to calculate H@k and MRR, which are defined as:

$$H@k = \frac{1}{N} \sum_i^N I(\text{rank}(i) < k), \quad (34)$$

$$MRR = \frac{1}{N} \sum_i^N \frac{1}{\text{rank}(i)}, \quad (35)$$

where N is the number of total samples, $\text{rank}(i)$ means the rank of the i -th ground truth entity in the

No.	Variant	H@1	H@3	H@5	MRR
-	OTMEL(OT)	83.30	92.39	94.83	88.27
F1	repl. ATT	80.60	91.66	94.50	86.70
F2	w/o FusM	78.94	90.76	93.60	85.42
M1	repl. ATT	81.72	92.34	95.06	87.61
M2	w/o UniM	73.95	88.12	91.86	81.85
A1	ATT(w/o OT)	78.33	91.16	94.33	85.28
O1	SoftPool \rightarrow Mean	80.80	90.90	93.91	86.55
O2	SoftPool \rightarrow Max	82.28	92.35	94.58	87.25
-	OT-MEL(KL)	82.57	92.06	94.44	87.78
K1	KD \rightarrow MSE	81.95	91.66	94.22	87.34
K2	KD \rightarrow UniKL	81.84	91.77	94.47	87.33
K3	Att(w/o KD)	78.33	91.16	94.33	85.28

Table 6: Results (%) of variants on Richpediamel data.

rank list of KG entities, $I(\cdot)$ stands for indicator function which is 1 if the subsequent condition is satisfied otherwise 0.

E Details of Variant Analysis

Table 6 displays all metrics from the variant experiments on the Richpediamel dataset.

F Analysis on Varying Token Counts

We also conducted extensive experiments on samples with different token counts in the Richpediamel dataset. The experimental results are presented in Figure 6. We categorized texts based on the number of tokens: texts with more than 20 tokens were classified as long texts, while those with fewer than 20 tokens were classified as short texts. Short texts contain fewer linking signals, making the short text group more challenging than the long text group. From our observations, we can draw the following conclusions: *Our model demonstrates stronger robustness in challenging scenarios compared to attention mechanism and baseline models.* In long text scenarios, our model slightly outperforms baseline models and attention mechanisms. However, in the more challenging short text scenarios, while our model only experiences a slight decrease in performance, the baseline models and attention mechanisms show a significant decline.

G Dissusion of the Generative Entity Linking

Generative Entity Linking (GEL) (De Cao et al., 2020) utilizes generative models (Lewis et al., 2019; Touvron et al., 2023) to directly generate the target entity of the mention. Due to the end-to-end

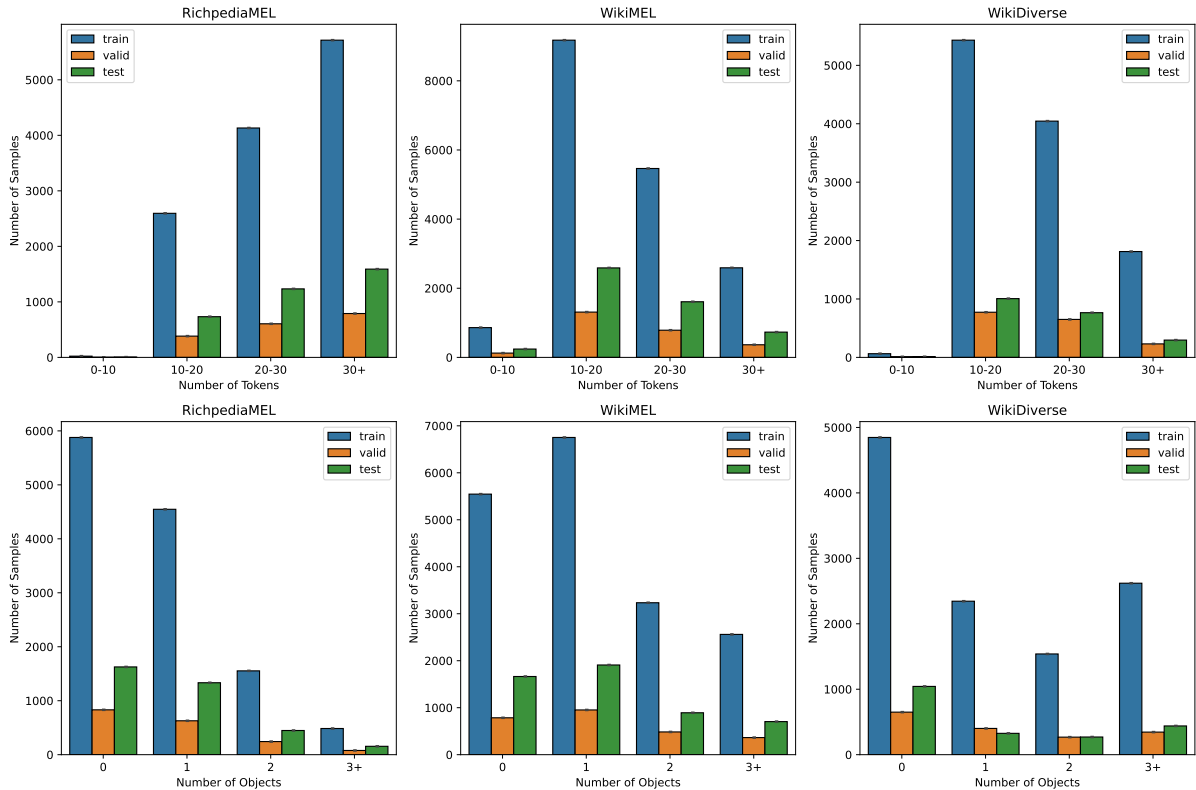


Figure 7: Statistics on text tokens and image objects across three datasets.

paradigm and competitive results, it has received extensive attention (Shi et al., 2023; De Cao et al., 2022; Huang et al., 2022; Yuan et al., 2022).

Generative methods use the constrained decoding mechanism to ensure that the entity names generated by the model must exist in the knowledge graph. Specifically, they construct a Trie of alternative entities, where each leaf in the tree is an entity name. However, they can hardly deal with entities with the same name, which we term as **lexical ambiguity problem** (Fortuny and Payrató, 2024). For example, *computer scientist Michael Jordan* and *basketball player Michael Jordan* share the same leaf node of the Trie, and we find that there are 14 pages (of entities) in Wikipedia called the same name *Michael Jordan*.

Model	Wikimel(Hit@1)	Wikidiverse(Hit@1)
GEMEL	82.6	86.3
OT-MEL(ours)	88.97	66.07

Table 7: Comparison with Generative Model.

To further investigate the lexical ambiguity problem, we count the number of these entities with the same name in the KG used for the three MEL datasets in the Table 5. We find that the Wikidi-

verse dataset has more data with lexical ambiguity problem (including 14,827 entities), and the Wikimel dataset has fewer data of that (including 616 entities). We further compare the results to the typical multimodal generative GEMEL (Shi et al., 2023), and show the results in Table 7. The results demonstrate that, on the *low lexical ambiguity* dataset Wikimel, GEMEL is lower than ours, and better than ours on the *high lexical ambiguity* dataset Wikidiverse. Although GEMEL performs better on the high lexical ambiguity dataset, it neglects the lexical ambiguity problem, thus cannot distinguish the entities with the same name.

Besides, GENER (De Cao et al., 2020), the unimodal generative method of the EL, can also hardly deal with this problem. It tries to solve the lexical ambiguity problem by generating distinguishable entity ids, but the final result drops by about 20% (De Cao et al., 2020). Considering the entity ids contain little semantics, how to generate correct ids is challenging for generative models.

With the emergence of LLMs, the GEL is a promising paradigms nowadays. However, how to solve the lexical ambiguity problem of GEL is one of the most important problems in its application, and we will explore it in the future work.