

Enhancing Multimodal Entity Linking with Jaccard Distance-based Conditional Contrastive Learning and Contextual Visual Augmentation

Cong-Duy Nguyen¹, Xiaobao Wu^{1*}, Thong Nguyen², Shuai Zhao¹,
Khoi Le³, Viet-Anh Nguyen¹, Yichao Feng¹, Anh Tuan Luu^{1*}

¹Nanyang Technological University, Singapore

²National University of Singapore, Singapore ³VinAI Research, Vietnam

{nguyentr003, xiaobao.wu, shuai.zhao, yichao.feng, anhtuan.luu}@ntu.edu.sg
thongnguyen050999@gmail.com

Abstract

Previous research on multimodal entity linking (MEL) has primarily employed contrastive learning as the primary objective. However, using the rest of the batch as negative samples without careful consideration, these studies risk leveraging “easy” features and potentially overlook essential details that make entities unique. In this work, we propose JD-CCL (Jaccard Distance-based Conditional Contrastive Learning), a novel approach designed to enhance the ability to match multimodal entity linking models. JD-CCL leverages meta-information to select negative samples with similar attributes, making the linking task more challenging and robust. Additionally, to address the limitations caused by the variations within the visual modality among mentions and entities, we introduce a novel method, CVaCPT (Contextual Visual-aid Controllable Patch Transform). It enhances visual representations by incorporating multi-view synthetic images and contextual textual representations to scale and shift patch representations. Experimental results on benchmark MEL datasets demonstrate the strong effectiveness of our approach.

1 Introduction

Multimodal Entity Linking (MEL) (Song et al., 2023; Yang et al., 2023; Shi et al., 2024) aims to connect visual and textual references to the entities in a multimodal knowledge graph (Song et al., 2023; Yang et al., 2023; Shi et al., 2024). As exemplified in Figure 3, MEL connects the Spider-Man references in the image and input sentence to the Spider-Man entity in the multimodal knowledge graph. This alleviates textual ambiguities by incorporating visual cues such as character portraits, which benefits various downstream applications, like information extraction (Yaghoobzadeh et al., 2017), question answering (Longpre et al., 2021), and search engines (Gerritse et al., 2022).

*Corresponding Authors.

Input Mention	Ground truth entity
Courtroom of the Supreme Court, 2015. Mention: Supreme Court	Supreme Court of the United Kingdom
	
Other entities in Knowledge base	
Supreme Court of Canada	United States Supreme Court Building
	

Figure 1: Example from WikiDiverse: The first two samples are the sentence input and the ground truth entity for its mention. The last two are other “Supreme Court” entities taken from the knowledge base.

However, current MEL methods are limited by their random negative sampling and understudy the potential issue of varied visual representation. First, these methods tend to employ contrastive learning to align input sentences and entities, but limited by its random negative sampling (Wang et al., 2022b; Luo et al., 2023). They generally randomly draw negative samples from the datasets. This may lead to unfair predictions due to numerous similar entities in the knowledge base, for example, different but highly similar “Supreme Court” entities in 1. Using them as negative samples hinders the connection between input sentences and correct entities. Second, existing MEL approaches overlooked the dissimilar visual modality between mention and target entity. For instance, 2 shows that the input




Figure 2: Illustration of challenge involving disparities within the visual modality. A sample from WikiDiverse includes the mention sentence and its image talking about the actor Doohan. Input sentence contains multiple entities from the knowledge base: Doohan, Star Trek, Scotty, Wikipedia.

image refers to an actor “Doohan”, but the input sentence includes four distinct entities. As a result, this greatly disturbs linking the input image to the entities in the knowledge base. For further specificity, the "Doohan" image on the left illustrates the sentence, offering significant utility to the model for user queries related to "Doohan" within the knowledge base due to image-text similarity. However, it could potentially introduce interference if the user intends to query "Wikipedia". Hence, pre-linking control of the visual modality, contingent upon the mentioned target, emerges as a crucial consideration in the Multimodal Entity Linking (MEL) task, particularly when other modalities co-exist alongside textual input.

To address the limitation of random negative sampling, we present a Jaccard Distance-based conditional contrastive learning (JD-CCL) approach for MEL. It leverages meta-information to mitigate the influence of prominent attributes during contrastive pre-training. JD-CCL automatically utilizes the information provided by each entity in the knowledge base to sample hard negative entities with nearly identical attributes. This strategy makes it more challenging for the model to link the mention with the correct entity, as it cannot rely on simple attributes (such as human, building, or animal) but must distinguish positive pairs from negative ones based on more complex attributes. To identify similar entities, we define a pre-processing stage that sorts the Jaccard similarity scores between the meta-attributes between entities in knowledge base. This process brings more difficult and nuanced comparisons during training, which improves the linking ability by diminishing the impact of “easy” attributes in contrastive training.

Moreover, to address the discrepancy challenge of visual modality in image-text pairs, we introduce the Contextual Visual-Aid Controllable Patch Transform (CVa-CPT) module. By leveraging synthetic images from a text-to-image diffusion, the model can better control essential features of the input image based on specific mention. Consequently, the model can selectively control useful patch representations from the input image rather than using all image patches with the same contribution. Additionally, for sentences containing multiple entities, each entity will have a unique input image, enhancing diverse visual representation according to the entity being queried.

Our contributions can be summarized as follows:

- We propose a novel Jaccard Distance-based Conditional Contrastive Learning (JD-CCL). This method leverages meta-attributes to draw conditional negative samples with similar features during training, enabling the model to identify exact characteristics for decision-making.
- We introduce the Contextual Visual-Aid Controllable Patch Transform (CVa-CPT) module, which enhances visual modality representations with synthetic images and contextual information from the mention sentence.
- We evaluate our approach on three benchmark datasets: WikiDiverse, RichpediaMEL, and WikiMEL, and show that our method outperforms previous state-of-the-art approaches.

2 Related Work

Multimodal Entity Linking In the landscape of social media and news, where content frequently

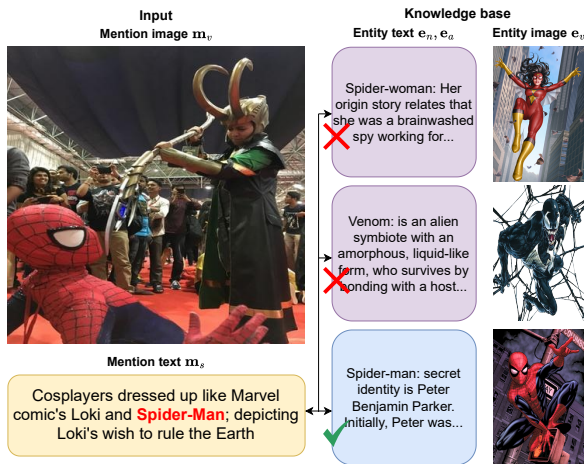


Figure 3: Task definition of Multimodal Entity Linking.

encompasses both text and visuals, integrating textual and visual information for entity linking is essential and pragmatic. In an early pioneering effort, (Moon et al., 2018) utilized images to enhance entity linking in response to ambiguous and incomplete references in social media posts. Advancing this, (Wang et al., 2022a) investigated the inter-modal interactions through a cross-attention mechanism between text and vision, integrating a gated hierarchical architecture. To mitigate the impact of noisy and irrelevant images, (Zhang et al., 2021) assessed the relevance of images by correlating their category with the semantic content of text mentions, filtering images using a predefined threshold. (Gan et al., 2021) compiled a dataset featuring extensive movie reviews, associated entities, and images. More recently, research (Zheng et al., 2022) has included the use of scene graphs from images to achieve object-level encoding, thus enhancing the granularity of visual semantic cues. MIMIC (Luo et al., 2023) presents a multi-grained multimodal interaction network designed for the multimodal entity linking task. GMEL (Shi et al., 2024) offers a Generative Multimodal Entity Linking framework employing LLMs that directly produces names of target entities. MMELL (Yang et al., 2023) introduces a module that jointly extracts features to learn representations from both visual and textual inputs for context and entity candidates. DWE (Song et al., 2023) describes a dual-way enhancement framework where the query benefits from refined multimodal information, and Wikipedia is used to augment the semantic representation of entities.

Contrastive learning Contrastive Learning (Chopra et al., 2005; Nguyen and Luu, 2021;

Nguyen et al., 2025, 2024a,b, 2023b) has become increasingly prominent in various fields (Nguyen et al., 2022, 2023a, 2024c; Wu et al., 2020b, 2022, 2023a,b, 2024a,c,b; Wei et al., 2024). Established training methods such as N-Pair Loss (Sohn, 2016), Triplet Margin Loss (Balntas et al., 2016), and ArcCon (Zhang et al., 2022) are foundational in metric learning. In supervised learning scenarios, key methodologies include Center loss (Wen et al., 2016), SphereFace (Liu et al., 2017), CosFace (Wang et al., 2018), and ArcFace (Deng et al., 2018), which are extensively applied in computer vision and natural language processing. Recent developments in contrastive learning incorporate additional conditional variables to enhance representation quality, involving auxiliary attributes (Tsai et al., 2021), information pertinent to the downstream task (Tian et al., 2020), downstream labels (Khosla et al., 2020; Kang and Park, 2020), or data embeddings (Tsai et al., 2022; Wu et al., 2020a). These advancements allow for the extension of contrastive self-supervised learning into weakly supervised (Tsai et al., 2021), semi-supervised (Tian et al., 2020), or fully supervised frameworks (Khosla et al., 2020). The work by (Khosla et al., 2020) specifically seeks to enhance fairness, rather than solely focusing on representation quality, by employing sensitive attributes from the dataset as the conditional variables in a self-supervised context.

3 Problem Formulation

A multimodal knowledge base includes a set of entities $E = \{\mathbf{E}_i\}_{i=1}^N$. Each entity is denoted as $\mathbf{E}_i = (e_{n_i}, e_{v_i}, e_{d_i}, e_{a_i})$, represent entity name, entity images, entity description, and entity attributes, respectively. In this representation, the \mathbf{E}_i components correspond to the entity's name, images, description, and attributes, respectively. As our study emphasizes linking local-level entities, textual input is formatted as sentences rather than as documents. Specifically, a mention and its context are represented as $\mathbf{M}_j = (\mathbf{m}_{w_j}, \mathbf{m}_{s_j}, \mathbf{m}_{v_j})$, where \mathbf{m}_{w_j} , \mathbf{m}_{s_j} , and \mathbf{m}_{v_j} denote the words of the mention, the sentence containing the mention and the associated image, respectively. The related entity of the mention \mathbf{M}_j in the knowledge base is \mathbf{E}_i . The MEL task is illustrated in Figure 3.

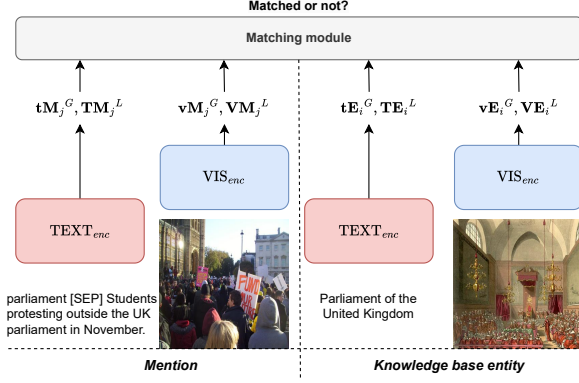


Figure 4: The overall architecture of MEL problem.

4 Methodology

In this section, we first introduce the overall framework and then go through the details of our approaches. We visualize the overall architecture in Figure 4.

4.1 Overall Architecture

Visual Encoding Following previous work, we utilize the pre-trained Vision Transformer (ViT) (Dosovitskiy et al., 2021) as our visual backbone to encode image features, denoted as VIS_{enc} . Each entity image e_{v_i} (or mention image m_{v_j}) is resized to $C \times H \times W$ pixels and divided into flattened 2D patches of size $n = H \times W/P^2$. These patches undergo projection and multi-layer transformer processes. Subsequently, a fully connected layer adjusts the output hidden states. The resulting hidden states, denoted as VE_i , include a global feature vE_i^G represented by the special token [CLS], and local features VE_i^L (similarly for mention vM_j^G and VM_j^L).

Textual Encoding We employ a Transformer-based architecture as the textual encoder (BERT (Devlin et al., 2019)), denoted as $TEXT_{enc}$. For mentions M_j , we concatenate the mention words with the sentence containing the mention to create the input sequence. We obtain tM_j^G and TM_j^L as the global and local textual features of mention $M_j \in \mathbb{R}^{(l_e+1) \times d_t}$, respectively. Here, d_T is the dimension of textual output features, and l_e is the length. For entities E_i , entity inputs are formed by concatenating the entity name with its attributes, denoted as e_{a_i} . Following the mentioned processing, the resulting hidden states TE_i are structured as $[t[CLS]^0; tE_i^1; \dots; tE_i^{l_e}] \in \mathbb{R}^{(l_e+1) \times d_t}$.

Matcher In this step, the model calculates the similarity matching score between input mention text-image pairs and knowledge entities to achieve correct matching. For instance, using the CLIP (Radford et al., 2021) template, we calculate the cosine similarity score between the global representation ([CLS] token) the mention input with the entity. Alternatively, MIMIC (Luo et al., 2023) introduced the Multi-Grained Multimodal Interaction Network, which captures and integrates the fine-grained representation of global and local cues of each modality. Overall, the matcher acts as a scoring function $\hat{y} = F(M, E)$ between the mention M and the entity E .

Objective Based on the score calculated above, we jointly train both the encoding layer and the matcher using a contrastive training loss function. This approach allows the model to learn to assign higher scores to positive mention-entity pairs and lower scores to negative mention-entity pairs. This loss function can be formulated as follows.

$$L = -\log \frac{\exp(F(M, E_{pos}))}{\exp(F(M, E_{pos})) + \sum_i \exp(F(M, E_{neg,i}))} \quad (1)$$

4.2 Conditional Contrastive learning

Conditional Contrastive Learning (CCL) (Ma et al., 2022) was proposed to enhance the fairness of contrastive self-supervised learning (SSL) methods. CCL mitigates the influence of sensitive attributes during contrastive pre-training by sampling positive and negative pairs that share the same attribute (e.g., gender, race). The CCL approach reduces the impact of the sensitive attribute Z by treating Z as the conditional variable between X and Y , focusing on scenarios where Z is readily available in the dataset (e.g., gender, race, or age). The proposed **Conditional Contrastive Learning** objective is designed to address this effectively.

$$L_{CCL} = \sup_f \mathbb{E}_{z \sim P_Z} \left[\mathbb{E}_{(x_i, y_i) \sim P_{X, Y|z}} \left[\frac{1}{n} \sum_{i=1}^n \log \frac{e^{f(x_i, y_i)}}{\frac{1}{n} \sum_{j=1}^n e^{f(x_i, y_j)}} \right] \right] \quad (2)$$

where the positive pairs $\{(x_i, y_i)\}_{i=1}^n$ represent samples drawn from the *conditional* joint distribution: $(x_i, y_i) \sim P_{X, Y|z}$, while the negative pairs $\{(x_i, y_{j \neq i})\}$ represent samples drawn from the product of conditional marginal distributions:

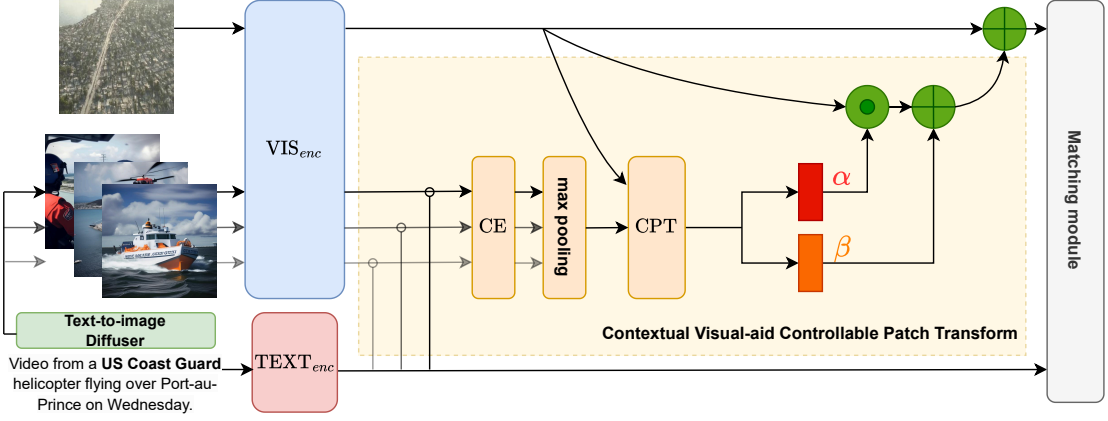


Figure 5: The overall architecture of Contextual Visual-aid Controllable Patch Transformation. The model incorporates a text-to-image diffuser, visual encoder, and text encoder, followed by a contextual visual-aid controllable patch transform (CVaCPT). The generated images and input sentence will help the model selectively control visual representation through the CVaCPT, including Contextual encoder, Max pooling, and Controllable patch transform.

$(x_i, y_{j \neq i}) \sim P_{X|z}P_{Y|z}$. The proposed Conditional Contrastive Learning (CCL) differs from conventional contrastive SSL by selecting positive and negative pairs from distributions conditioned on a sensitive attribute referred to as Z . The CCL approach reduces the influence of the sensitive attribute Z by treating Z as the conditional variable between X and Y . We focus on scenarios where Z is readily available in the dataset (e.g., gender, race, or age).

Conditional sampling We first define a preprocessing stage to identify hard negative samples before the training phase. Each entity \mathbf{E}_i in the knowledge base contains a set of predefined attributes. We calculate the Jaccard distance between two entities i and j as the similarity score between the two entities:

$$J(i, j) = \frac{|\mathbf{e}_{a_i} \cap \mathbf{e}_{a_j}|}{|\mathbf{e}_{a_i} \cup \mathbf{e}_{a_j}|} \quad (3)$$

For each entity \mathbf{E}_i , we calculate the total Jaccard distance with the knowledge base $J_i = [J(i, 1), J(i, 2), \dots, J(i, j)]$, where $1 \leq j \leq n_e$ and $j \neq i$. We store these similarity scores and sort them. Then, we select the top k scores as the k most similar or hard negative samples of \mathbf{E}_i : $\hat{J}_i = \text{top}_k, J_i$.

Jaccard Distance-based CCL After determining the top k negative samples, we calculate the conditional contrastive learning meta-attribute JD-CCL:

$$\begin{aligned} \mathcal{L}_{JD-CCL} = & E_{(\mathbf{M}, \mathbf{E}_{pos}) \sim P_{\mathbf{M}, \mathbf{E}}, z \sim \hat{J}_{\mathbf{E}_{pos}}, \{\mathbf{E}_{neg, j}\}_{j=1}^k \sim P_{\mathbf{E}|Z=z_i}} \\ & \left[\log \frac{e^{F(\mathbf{M}, \mathbf{E}_{pos})}}{e^{F(\mathbf{M}, \mathbf{E}_{pos})} + \sum_{j=1}^k e^{F(\mathbf{M}, \mathbf{E}_{neg, j})}} \right] \quad (4) \end{aligned}$$

where positive pairs $(\mathbf{M}_i, \mathbf{E}_{pos, i})$, consisting of a mention input and its linked entity in the knowledge base, are drawn randomly from the dataset. After sampling a positive pair, negative entities $\{\mathbf{E}_{neg, j}\}_{j=1}^k$ are drawn from the set \hat{J}_i of the entity $\mathbf{E}_{pos, i}$. Unlike CCL for fair representation, this approach samples negative entities based on their meta-information, thereby covering different attributes, whereas CCL is based on predefined conditions Z such as gender and race. Furthermore, the difference between JD-CCL and Contrastive Learning can be described as follows: JD-CCL samples negative entities from conditions $z \in Z$, resulting in negative pairs that possess more general features. This increased diversity in features makes the model more challenging to train.

4.3 Contextual Visual-aid Controllable Patch Transformation

To address the issue of dissimilarity between visual information in mention inputs and knowledge entities, we propose the Contextual Visual-aid Controllable Patch Transformation (CVaCPT) module. CVaCPT adaptively influences the mention visual input by applying a transformation to each patch representation based on prior information. Technically, we utilize the Stable Diffusion model to

generate an image s_j corresponding to the input text x_j . Here, m_j is the mention sentence prefixed with "An image of". Thus, the same sentence with different mentions we want to match will create different images of the target. Using a visual encoder VIS_{enc} , we obtain a global representation VEs_j^G of the synthetic image s_j .

Controllable patch transformation (CFT) The generated image with the target reference text in the sentence helps the model selectively control the important features of the mention input. We propose the Controllable Feature Transformation (CFT) module to manage the information flow from the visual encoder, VE_j^G and VE_j^L , before feeding it into the matcher. Specifically, as shown in Fig. 5, the synthetic features VEs_j^G combined with the text mention representation TE_j^G are used to fine-tune the visual features of the mention, VE_j^G and VE_j^L , through the transformation of features using the affine parameters α and β :

$$\hat{VE}_j^G = VE_j^G + w(\alpha^G \odot VE_j^G + \beta^G) \quad (5)$$

$$\hat{VE}_j^L = VE_j^L + w(\alpha^L \odot VE_j^L + \beta^L) \quad (6)$$

$$\alpha^G, \beta^G = A^G([VE_j^G \cdot VTEs_j]) \quad (7)$$

$$\alpha^L, \beta^L = A^L([VE_j^L \cdot VTEs_j]) \quad (8)$$

$$VTEs_j = CE([VEs_j^G \cdot TE_j^G]) \quad (9)$$

where P_{θ^X} denotes a stack of convolutions that predicts α^X and β^X from the concatenated (\cdot) features $[VE_j \cdot CE([VEs_j^G \cdot TE_j^G])]$. CE is the contextual visual-aid module, which receives a concatenation of the global representation of the synthetic image VEs_j^G and the global representation of the input mention TE_j^G . An adjustable coefficient $w \in [0, 1]$ is then used to control the relative importance of the inputs. With the CPT module, the mention image representations are adjusted to match the mentioned entity in the sentence, removing noisy and unhelpful patches needed to match in the knowledge base.

Multiview Synthetic image We found that images generated from the diffuser model can include noisy features (e.g., background details not related to the mention). Thus, using multiple images can help the CFT module select important and helpful features from synthetic images. We generate n_s different synthetic images from the same prompt $s_{j,h}$, where $h \in \{1, \dots, n_s\}$. Each global visual repre-

sentation is concatenated with the global textual representation before being fed into CE:

$$V\hat{TE}s_j = [VTEs_{j,1}, \dots, VTEs_{j,n_s}] \quad (10)$$

$$VTEs_{j,h} = CE([VEs_{j,h}^G \cdot TE_j^G]) \quad (11)$$

We apply max-pooling to calculate the maximum value for patches of a feature map across multiple images:

$$VTEs_j = \max_pooling(V\hat{TE}s_j) \quad (12)$$

This max_pooling operation emphasizes the presence of strong/important features and diminishes the noise.

5 Experiments setups

5.1 Datasets

In the experiments, we adopt three benchmark MEL datasets: **(i) WikiDiverse** (Wang et al., 2022b) is a meticulously curated MEL dataset featuring the appendixverse array of contextual topics and entity types sourced from Wikinews. Comprising 8,000 image-caption pairs, it employs Wikipedia as its underlying knowledge base. **(ii) WikiMEL** (Wang et al., 2022a) comprises over 22,000 multimodal sentences gathered from Wikipedia entity pages. **(iii) RichpediaMEL** (Wang et al., 2020) Initially, entities were extracted from Richpedia for the creation, followed by the acquisition of multimodal data from Wikidata (Vrandečić and Krötzsch, 2014). Individuals constitute the primary entity category in both WikiMEL and RichpediaMEL.

5.2 Baselines

We compared our method with various competitive VLP models: **CLIP** (Radford et al., 2021), **ViLT** (Kim et al., 2021), **ALBEF** (Li et al., 2021), **METER** (Dou et al., 2022), **MIMIC** (Luo et al., 2023). Detailed descriptions are provided in the Appendix A.

5.3 Metrics

We calculated the similarity between a mention and all entities in the knowledge base (KB) to measure their aligning probability. The similarity scores are sorted in descending order to calculate **H@1**, **H@3**, **H@5**, and **MRR**. **H@k** indicates the hit rate of the ground truth entity when only considering the top-k ranked entities: **H@1,3,5** represent the

Model	WikiDiverse				RichpediaMEL				WikiMEL			
	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR	H@1	H@3	H@5	MRR
CLIP	61.21	79.63	85.18	71.69	67.78	85.22	90.04	77.57	83.23	92.10	94.51	88.23
ViLT	34.39	51.07	57.83	45.22	45.85	62.96	69.80	56.63	72.64	84.51	87.86	79.46
ALBEF	60.59	75.59	81.30	69.93	65.17	82.84	88.28	75.29	78.64	88.93	91.75	84.56
METER	53.14	70.93	77.59	63.71	63.96	82.24	87.08	74.15	72.46	84.41	88.17	79.49
MIMIC	63.51	81.04	86.43	73.44	81.02	91.77	94.38	86.95	87.98	95.07	96.37	91.82
Ours	70.25	83.78	88.06	78.09	83.38	91.91	94.18	88.14	89.28	95.31	96.84	92.62

Table 1: Performance comparison on three MEL tasks: WikiDiverse, RichpediaMEL, and WikiMEL. We trained the models using a random seed of 5, presenting the average of our findings. Our experiments included paired t-tests with $p < 0.05$, demonstrating overall significant improvement.

Setting	H@1	H@3	H@5	MRR
Our	70.25	83.78	88.06	78.09
w/o LD-CCL	66.55	82.57	87.53	75.72
w/o CVaCPT	67.80	81.85	86.62	76.04
(1) Number of negative samples in LD-CCL				
2	67.88	82.57	87.24	77.26
4	69.00	83.01	87.34	77.08
6	69.53	84.16	88.30	77.80
(2) Number of synthetic images in CVaCPT module				
1	66.68	82.85	87.10	76.04
2	68.47	83.78	88.16	77.03
3	68.67	83.54	86.86	76.87

Table 2: Ablation results on WikiDiverse.

hit rates in the top 1, 3, and 5, respectively. **MRR** represents the mean reciprocal rank of the ground truth entity: $\frac{1}{N} \sum_{i=1}^N \frac{1}{rank(i)}$.

5.4 Implementations

We initialize the CLIP model (Vit-Base-Patch32) as both our visual encoder VIS_{enc} and text encoder $TEXT_{enc}$, from HuggingFace. For visual input, all images (mention, knowledge entity and synthetic image) are rescaled to 224×224 . With the mention or entity input that do not have image, will initialize with blank image (while) with zero padding. For textual input, we use CLIP tokenizer, which constructs based on byte-level Byte-Pair Encoding. We implement linear neural networks to project the hidden representation from CLIP model to 96 and setup the matcher similar to MIMIC. Since our CVaCPT module and contrastive objective are flexibly integrated with various models, we built upon the implementation of the state-of-the-art MIMIC (Luo et al., 2023) by incorporating CVaCPT and replacing the original contrastive learning objective with our JD-CCL approach.

Model	H@1	H@3	H@5	MRR
WikiDM				
MIMIC	74.77	89.12	92.25	82.63
Our	78.77	90.25	92.92	85.16
WikiDR				
MIMIC	72.41	84.75	89.00	79.63
Our	77.32	88.65	91.70	83.59

Table 3: Ablation results WikiDM and WikiDR.


6 Experimental results

Following the methodology of previous studies, we conducted our model five times using the same hyperparameter settings. The average results are presented in Table 1. An overview of our results reveals that our model consistently matches or surpasses the benchmarks set by baseline approaches. Our proposed method shows significant improvements across various datasets. Notably, we achieved a substantial increase of 6.74% in accuracy for H@1 on WikiDiverse, and improvements of 2.36% and 1.30% on RipediaMEL and WikiMEL, respectively. For H@3, H@5, and MRR, we also achieved better performance on most datasets. To be more specific, we outperform the state-of-the-art (SOTA) by 2.74%, 1.63%, and 4.65% on WikiDiverse, and achieve slight improvements on the other two datasets, except for H@5 metric on WikiMEL.


7 Analysis

7.1 Regularization

To validate the significance of each of our objectives, we conducted additional experiments where one of our proposed methods was removed. As illustrated in Table 2, the results are presented in the first three rows, indicating that incorporating all the proposed approaches leads to the best performance. The second row shows that not applying



(a) Sentence: Barack Obama, Michelle Obama, Jill Biden and Joe Biden at the United States Vice Presidential announcement on August 23, 2008. Mention: Barack Obama.



(b) Sentence: The exterior of the Modern One building of Edinburgh’s Scottish National Gallery of Modern Art. Mention: Modern One.

Figure 6: Visualization of CVaCPT module

JD conditional contrastive learning harms performance. This result can be attributed to the fact that, during training, hard negative samples play an important role in this task. In the third row of the table, we demonstrate that the removal of CVaCPT significantly affects performance across all metrics.

7.2 Number of negative samples and number of synthetic images

We conducted an experiment to better understand the impact of the number of negative samples, k , in JD-CCL and the number of synthetic images, s , in CVaCPT. We set k to values in the set 2, 4, 6 and s to values in the set 1, 2, 3. As shown in the first three rows of Table 2, a larger number of condi-

tional negative samples ($k = 4, 6$) achieves better results compared to a smaller number ($k = 2$), although there is only a small difference between $k = 4$ and $k = 6$. Regarding the number of synthetic images, using $s = 1$ means not applying the max-pooling operation. As shown in the last three rows, multiple views of visual augmentation outperform a single view ($s = 1$) with only one synthetic image. However, increasing the number of images from 2 to 3 does not significantly improve the model.

7.3 Larger knowledge base evaluation

While the benchmarks WikiMEL and Richpedi-aMEL achieved good performance, WikiDiverse appears to be a more challenging dataset. To enhance the ability to handle a larger knowledge base, we merged the two easier benchmarks with the more challenging WikiDiverse to create two larger datasets (double in size): WikiDM (WikiDiverse + WikiMEL) and WikiDR (WikiDiverse + Richpedi-aMEL). The experimental results are presented in Table 3. As shown, our approach outperforms MIMIC across all metrics on both datasets. Specifically, our method achieves an improvement of 4.00% and 4.91% in the H@1 metric on WikiDM and WikiDR, respectively, and outperforms in all other metrics as well.

7.4 Visualization

In this appendix, we provide a visualization of the Controllable Patch Transformation module, as illustrated in Figures 6a and 6b. We visualize the calculation of alpha and beta for patches in the image corresponding to the input mention (in each caption) and synthetic images. As shown in Figure 6a, the model focuses on people in the image (alpha and beta values are large) while reducing the impact of the background (sky, building) (alpha and beta values are small). In contrast, Figure 6b shows a building covering the entire image, where all alpha and beta values are equally high, indicating that all image patches are informative for the matcher.

8 Conclusion

This paper introduces JD-CCL, a novel conditional contrastive learning approach based on Jaccard Distance. This method leverages meta-attributes to draw conditional negative samples with similar features during training, allowing the model to identify specific characteristics for prediction accurately.

Additionally, our proposed Contextual Visual-Aid Controllable Patch Transform (CVa-CPT) module enhances visual modality representation by utilizing synthetic images and contextual information from the mention sentence. Our framework paves the way for further exploration in this domain, promoting multimodal and knowledge base application advancements.

Limitations

Despite the improvements our approach achieves in the MEL task, JD-CCL and CVaCPT still have the following limitations. CVaCPT uses a Text-to-Image diffuser model to generate synthetic images, which requires additional time and computational resources. Furthermore, using synthetic images from a different domain might introduce some noise due to the domain gap. Secondly, JD-CCL includes a condition sampling stage, which is a preprocessing step to determine the negative samples, and this stage also takes time. Additionally, we automatically calculate the Jaccard distance between the attribute sets of two entities. If the dataset is not well-annotated, this could negatively impact the training stage.

Acknowledgement

This research is supported by DSO Singapore under the research grant DSOCL23216.

References

- Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. 2016. [Learning local feature descriptors with triplets and shallow convolutional neural networks](#). In *British Machine Vision Conference*.
- S. Chopra, R. Hadsell, and Y. LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546 vol. 1.
- Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. 2018. [Arcface: Additive angular margin loss for deep face recognition](#). *CoRR*, abs/1801.07698.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*. OpenReview.net.
- Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chenguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, Zicheng Liu, and Michael Zeng. 2022. An empirical study of training end-to-end vision-and-language transformers. In *CVPR*, pages 18145–18155. IEEE.
- Jingru Gan, Jinchang Luo, Haiwei Wang, Shuhui Wang, Wei He, and Qingming Huang. 2021. Multimodal entity linking: A new dataset and A baseline. In *ACM Multimedia*, pages 993–1001. ACM.
- Emma J. Gerritse, Faegheh Hasibi, and Arjen P. de Vries. 2022. Entity-aware transformers for entity search. In *SIGIR*, pages 1455–1465. ACM.
- Minguk Kang and Jaesik Park. 2020. Contragan: Contrastive learning for conditional image generation. *Advances in Neural Information Processing Systems*, 33:21357–21369.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*.
- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, pages 9694–9705.
- Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. 2017. [Sphereface: Deep hypersphere embedding for face recognition](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6738–6746.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-based knowledge conflicts in question answering. In *EMNLP (1)*, pages 7052–7063. 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](#). *Preprint*, arXiv:2307.09721.
- Martin Q. Ma, Yao-Hung Hubert Tsai, Paul Pu Liang, Han Zhao, Kun Zhang, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022. [Conditional contrastive learning for improving fairness in self-supervised learning](#). *Preprint*, arXiv:2106.02866.

- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. 2018. Multimodal named entity disambiguation for noisy social media posts. In *ACL (1)*, pages 2000–2008. Association for Computational Linguistics.
- Cong-Duy Nguyen, Thong Nguyen, Duc Anh Vu, and Luu Anh Tuan. 2023a. [Improving multimodal sentiment analysis: Supervised angular margin-based contrastive learning for enhanced fusion representation](#). *Preprint*, arXiv:2312.02227.
- Thong Nguyen, Yi Bin, Xiaobao Wu, Xinshuai Dong, Zhiyuan Hu, Khoi Le, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2025. Meta-optimized angular margin contrastive framework for video-language representation learning. In *European Conference on Computer Vision*, pages 77–98. Springer.
- Thong Nguyen and Anh Tuan Luu. 2021. [Contrastive learning for neural topic model](#). *Advances in Neural Information Processing Systems*, 34.
- Thong Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy Nguyen, See-Kiong Ng, and Luu Anh Tuan. 2023b. Demaformer: Damped exponential moving average transformer with energy-based modeling for temporal language grounding. *arXiv preprint arXiv:2312.02549*.
- Thong Nguyen, Xiaobao Wu, Anh Tuan Luu, Zhen Hai, and Lidong Bing. 2022. [Adaptive contrastive learning on multimodal transformer for review helpfulness prediction](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10085–10096, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Thong Thanh Nguyen, Yi Bin, Xiaobao Wu, Zhiyuan Hu, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024a. Multi-scale contrastive learning for video temporal grounding. *arXiv preprint arXiv:2412.07157*.
- Thong Thanh Nguyen, Xiaobao Wu, Yi Bin, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024b. Motion-aware contrastive learning for temporal panoptic scene graph generation. *arXiv preprint arXiv:2412.07160*.
- Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024c. [Topic modeling as multi-objective optimization with setwise contrastive learning](#). In *The Twelfth International Conference on Learning Representations*.
- 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 *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Senbao Shi, Zhenran Xu, Baotian Hu, and Min Zhang. 2024. [Generative multimodal entity linking](#). *Preprint*, arXiv:2306.12725.
- Kihyuk Sohn. 2016. [Improved deep metric learning with multi-class n-pair loss objective](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Shezheng Song, Shan Zhao, Chengyu Wang, Tianwei Yan, Shasha Li, Xiaoguang Mao, and Meng Wang. 2023. [A dual-way enhanced framework from text matching point of view for multimodal entity linking](#). *Preprint*, arXiv:2312.11816.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. 2020. What makes for good views for contrastive learning? *arXiv preprint arXiv:2005.10243*.
- Yao-Hung Hubert Tsai, Tianqin Li, Weixin Liu, Peiyuan Liao, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. Integrating auxiliary information in self-supervised learning. *arXiv preprint arXiv:2106.02869*.
- Yao-Hung Hubert Tsai, Tianqin Li, Martin Q Ma, Han Zhao, Kun Zhang, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2022. Conditional contrastive learning with kernel. *arXiv preprint arXiv:2202.05458*.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85.
- Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Zhifeng Li, Dihong Gong, Jingchao Zhou, and Wei Liu. 2018. [Cosface: Large margin cosine loss for deep face recognition](#). *CoRR*, abs/1801.09414.
- Meng Wang, Haofen Wang, Guilin Qi, and Qiushuo Zheng. 2020. Richpedia: A large-scale, comprehensive multi-modal knowledge graph. *Big Data Res.*, 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, SIGIR '22*, page 938–948, New York, NY, USA. Association for Computing Machinery.
- 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. In *ACL (1)*, pages 4785–4797. Association for Computational Linguistics.
- Jie Wei, Guanyu Hu, Xinyu Yang, Anh Tuan Luu, and Yizhuo Dong. 2024. Learning facial expression and body gesture visual information for video emotion recognition. *Expert Systems with Applications*, 237:121419.

- Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. 2016. A discriminative feature learning approach for deep face recognition. In *Computer Vision – ECCV 2016*, pages 499–515, Cham. Springer International Publishing.
- Mike Wu, Milan Mosse, Chengxu Zhuang, Daniel Yamins, and Noah Goodman. 2020a. Conditional negative sampling for contrastive learning of visual representations. *arXiv preprint arXiv:2010.02037*.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, Chaoqun Liu, Liang-Ming Pan, and Anh Tuan Luu. 2023a. [Infocfm: A mutual information maximization perspective of cross-lingual topic modeling](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13763–13771.
- Xiaobao Wu, Xinshuai Dong, Thong Nguyen, and Anh Tuan Luu. 2023b. [Effective neural topic modeling with embedding clustering regularization](#). In *International Conference on Machine Learning*. PMLR.
- Xiaobao Wu, Xinshuai Dong, Liangming Pan, Thong Nguyen, and Anh Tuan Luu. 2024a. [Modeling dynamic topics in chain-free fashion by evolution-tracking contrastive learning and unassociated word exclusion](#). In *Findings of the Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics.
- Xiaobao Wu, Chunping Li, Yan Zhu, and Yishu Miao. 2020b. [Short text topic modeling with topic distribution quantization and negative sampling decoder](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1772–1782, Online.
- Xiaobao Wu, Anh Tuan Luu, and Xinshuai Dong. 2022. [Mitigating data sparsity for short text topic modeling by topic-semantic contrastive learning](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2748–2760, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xiaobao Wu, Fengjun Pan, and Anh Tuan Luu. 2024b. [Towards the TopMost: A topic modeling system toolkit](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 31–41, Bangkok, Thailand. Association for Computational Linguistics.
- Xiaobao Wu, Fengjun Pan, Thong Nguyen, Yichao Feng, Chaoqun Liu, Cong-Duy Nguyen, and Anh Tuan Luu. 2024c. [On the affinity, rationality, and diversity of hierarchical topic modeling](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Yadollah Yaghoobzadeh, Heike Adel, and Hinrich Schütze. 2017. [Noise mitigation for neural entity typing and relation extraction](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1183–1194, Valencia, Spain. Association for Computational Linguistics.
- 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.
- Li Zhang, Zhixu Li, and Qiang Yang. 2021. Attention-based multimodal entity linking with high-quality images. In *DASFAA (2)*, volume 12682 of *Lecture Notes in Computer Science*, pages 533–548. Springer.
- Yuhao Zhang, Hongji Zhu, Yongliang Wang, Nan Xu, Xiaobo Li, and Binqiang Zhao. 2022. [A contrastive framework for learning sentence representations from pairwise and triple-wise perspective in angular space](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4892–4903, Dublin, Ireland. Association for Computational Linguistics.
- Qiushuo Zheng, Hao Wen, Meng Wang, and Guilin Qi. 2022. Visual entity linking via multi-modal learning. *Data Intell.*, 4(1):1–19.

CLIP	Normal feature	Pooling feature
Mention images	0.406	0.430
KB entity images	0.552	0.585
DinoV2	Normal feature	Pooling feature
Mention images	0.178	0.196
KB entity images	0.310	0.344

Table 4: Comparison of similarity scores between the mean of individual images and pooling features with KB entity images.





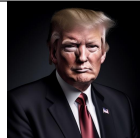

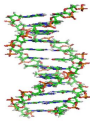





Mention image	KB entity image	Synthetic images			Pooling feature
Mention: President Trump holds a Bible in front of St. John’s Episcopal Church.					
KB entity: President of the United States from 2017 to 2021.					
					
	Dino	0.287	0.306	0.334	0.393
	CLIP	0.440	0.435	0.461	0.518
Mention: Radiotherapy and chemotherapy act together to damage the DNA of cancer cells. According to this new study, gemcitabine potentiates this effect.					
KB entity: Therapy using ionizing radiation					
					
	Dino	0.381	0.360	0.433	0.457
	CLIP	0.432	0.519	0.619	0.631

Table 5: Visualization of synthetic images in the Contextual Visual-aid Controllable Patch Transformation module.

A Baselines

We followed the baseline setup from (Luo et al., 2023):

CLIP leverages two Transformer-based encoders to generate visual and textual representations, pre-training on vast amounts of noisy web data using contrastive loss between 2 global representations of each modality.

METER adopts a co-attention schema to explore the semantic relationships between different modalities, with each layer comprising a self-attention module, a cross-attention module, and a feed-forward network.

ViLT applies deep modality interaction through a stack of Transformer layers and proposes utilizing early, shallow textual and visual embeddings.

ALBEF initially aligns visual and textual features through image-text contrastive loss and then integrates them using a multimodal Transformer encoder. It further employs momentum distillation to enhance learning from noisy data.

MIMIC represents the current state-of-the-art in MEL tasks across three datasets. MIMIC introduces three modules: the Text-based Global-Local Interaction Unit (TGLU), the Vision-based Dual Interaction Unit (VDLU), and the Cross-Modal Fusion-based Interaction Unit (CMFU), which are designed to capture and integrate fine-grained representations in abbreviated text and implicit visual cues.

	(1)	(2)	(3)	(4)
Wikidiverse				
Train	8752	2517	66	16
Valid	1275	384	5	0
Test	1599	461	15	3
RichpediaMEL				
Train	10992	86	1375	10
Valid	3194	19	348	1
Test	1545	16	219	0
WikiMEL				
Train	17725	367	0	0
Valid	2543	42	0	0
Test	5066	103	0	0

Table 6: (1) Both mention and knowledge base have the image. (2) Mention has an image but its knowledge base does not have an image. (3) Mention does not have an image but its knowledge base has an image. (4) Both the mention and knowledge base do not have the image.

B Noisy synthetic images

We conducted additional experiments to gain further insights into the behavior of the diffusion model in CVaCPT. For each sample in the dataset, we first used the visual encoder to extract features from the mention input image, the knowledge base entity image, and the synthetic image. We then calculated the cosine similarity between each synthetic image and the pooled synthetic image, comparing them to the mention and knowledge base entity images. The mean similarity scores are reported in Table 4. As shown, the similarity scores for pooled visual features are higher than those obtained from independent images. This indicates that using multiple synthetic images with a pooling operation helps mitigate the noise introduced by the diffusion model. We also visualize a selection of samples from the datasets in Table 5, illustrating the diversity of images generated by the model. The similarity scores exhibit a wide distribution across individual images, while the pooled features consistently provide better alignment.

C H3 and H5 insights

Our method demonstrated improved performance on MEL benchmarks compared to baselines. However, the gains in H3 and H5 metrics were marginal. To investigate this issue, we conducted additional experiments and data analysis. First, we trained and evaluated our method alongside MIMIC under identical settings, revealing that both models faced the same challenges when tested on the same samples in the test set. Upon further examination, we observed that while the correct entities were often ranked within the top 3 or 5, these mentions were consistently mislinked to entities with the same name in the knowledge base (i.e., same name but representing different entities). Table 7 illustrates two randomly selected failure cases, highlighting several factors that contributed to incorrect predictions: (1) the entities involved shared identical names, (2) these entities were never linked to any mention in the training set, and (3) many of the input entities or mentions lacked corresponding images. For issue (1), we provide a list of duplicate entity names along with their frequency larger than 10 in the knowledge base in Sec D. Regarding issue (2), the sheer size of the knowledge base relative to the inputs of mention in the data set resulted in numerous entities, particularly those with duplicate names, being excluded from the training process. Finally, as shown in Table 6, a portion of entity-mention pairs across the three datasets used in this study lacked images for at least one of the mentions or knowledge base entities. These challenges, (1), (2), and (3), collectively hinder the model’s ability to fully optimize for such cases.


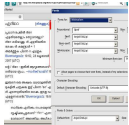


	Mention input	KB entity	Other KBs		
Image				Missing	Missing
Description	The Wikipedia home page as viewed on a Netscape browser.	Netscape - American computer services company.	Netscape Communicator - Discontinued Internet software suite	Netscape - family of web browsers	Netscape 6 - version 6 of the Netscape web browser
MIMIC		3	1	2	4
Our method		3	1	2	4
Image		Missing	Missing	Missing	Missing
Description	England did win this October 13, 2007 match versus Estonia , 3-0 at Wembley Stadium	Estonia sovereign state in northeastern Europe	Estonia men's national basketball team	Estonia national football team results	Estonia national football team
MIMIC		4	3	2	1
Our method		3	4	1	2

Table 7: Visualization of failed cases, samples taken from WikiDiverse

D Similar knowledge entities

Washington County (19), John Murphy (18), Greg Smith (18), Black Friday (17), St. John's Episcopal Church (17), Vanity Fair (17), Oliver Twist (17), Washington Township (17), Monroe County (17), Union County (17), Inside Out (16), Masquerade (16), Steve Smith (16), Lockdown (16), Wayne County (16), Montgomery County (16), Alice in Wonderland (15), Red Line (14), House of Representatives (14), Mount Vernon (14), Woodstock (13), Gordon Brown (13), Moby Dick (13), Paul Ryan (13), Government House (13), John Baird (13), Burlington (12), John Edwards (12), La Grange (12), Victoria Park (12), Lake County (12), Stephen Smith (12), Grafton (12), Michael Gallagher (12), Peter Marshall (12), Scarecrow (12), Marshall County (12), Independence Day (12), Greenfield (12), Polk County (12), Ceres (11), Arcadia (11), Homecoming (11), Camera Obscura (11), Equinox (11), Jack Thompson (11), Black Widow (11), San Fernando (11), Scott Brown (11), Supernova (11), White City (11), William Johnson (11), Brian Moore (11), Brush Creek (11), Chief of the Defence Staff (11), Calhoun County (11), Fayette County (11).