



Dynamic Uncertainty Learning with Noisy Correspondence for Text-Based Person Search

Haoming Ji^{1,+[0009-0008-9771-2413]} and Zequn Xie^{2,+[0009-0007-5021-7565]}

¹ Beijing University of Posts and Telecommunications, Beijing 100876, China
jihaoming@bupt.edu.cn

² Zhejiang University, Hangzhou, China
zqxie@zju.edu.cn

⁺ These authors contribute equally to this work

Abstract. Text-to-image person search aims to identify an individual based on a text description. To reduce data collection costs, large-scale text-image datasets are created from co-occurrence pairs found online. However, this can introduce noise, particularly mismatched pairs, which degrade retrieval performance. Existing methods often focus on negative samples, which amplify this noise. To address these issues, we propose the Dynamic Uncertainty and Relational Alignment (DURA) framework, which includes the Key Feature Selector (KFS) and a new loss function, Dynamic Softmax Hinge Loss (DSH-Loss). KFS captures and models noise uncertainty, improving retrieval reliability. The bidirectional evidence from cross-modal similarity is modeled as a Dirichlet distribution, enhancing adaptability to noisy data. DSH adjusts the difficulty of negative samples to improve robustness in noisy environments. Our experiments on three datasets show that the method offers strong noise resistance and improves retrieval performance in both low- and high-noise scenarios.

Keywords: Text-Based Person Search, Noisy Correspondences, Cross-modal Uncertainty Learning.

1 Introduction

Person Re-identification (Re-ID) is a key task in computer vision, important for applications like intelligent video surveillance, urban security, and smart retail [17]. Text-to-image person search aims to identify a person based on textual descriptions. However, current Re-ID methods rely on specific images of individuals, which may not be available in real-world emergency situations. In these cases, eyewitness descriptions can be the only information. To address this, text-based person search is used to identify individuals from image collections based on text queries.

Text-to-image person search is a sub-domain of both image-text retrieval [25, 31, 35] and image-based person Re-ID [7, 29, 36]. Text descriptions offer an intuitive way to express a person’s features, often easier to access than images.

This has led to growing interest in the field, with applications in personal photo searches and public safety. The main challenge is accurately measuring the similarity between text descriptions and images to retrieve the correct individual from the image database based on the provided text query.

To address these issues, current methods [2] use techniques to increase the similarity of positive text-image pairs while reducing that of negative pairs. Approaches like common representation learning [24] and similarity learning [10] have been applied. While these methods show promise, many depend on large, well-annotated datasets. To reduce data collection costs, co-occurring text-image pairs are often gathered from online sources, creating a large, cost-effective cross-modal dataset. However, such data often contains noise, known as noisy correspondence, which weakens the reliability of cross-modal relationships and affects retrieval performance. This problem is especially severe with hinge-based triplet ranking loss involving hard negatives [24, 10], where the hard learning approach becomes more vulnerable to noise, lowering retrieval accuracy, as shown in our experiments (e.g., Table 1).

Noisy correspondence is closely related to learning with noisy labels [30, 14], a concept studied in classification tasks. Various methods, like co-teaching [14] and robust loss functions [30], have been developed to address noisy labels. However, noisy correspondence involves misalignments between cross-modal pairs [16, 12, 26, 40], which differs from noisy labels in classification. Traditional robust learning techniques are inadequate for noisy correspondence, as it involves uncertainty at the instance level rather than the category level. As a result, noisy correspondence is much more complex than noisy labels, as the number of instances usually exceeds the number of categories by far.

To address these challenges, we propose the Dynamic Uncertainty and Relational Alignment (DURA) framework, designed for text-to-image person search in noisy environments. DURA introduces a new Cross-modal Evidential Learning (CEL) method to identify and handle uncertainty caused by noise, helping to isolate unreliable pairs and improve data reliability. To deal with unreliable hard negatives, the Dynamic Softmax Hinge (DSH) mechanism increases the difficulty of negative samples during training, reducing the impact of noisy correspondence. By combining CEL and DSH, DURA effectively distinguishes and uses both clean and noisy data. CEL models bidirectional evidence as a Dirichlet distribution based on cross-modal similarity, addressing uncertainty from noisy correspondence. This evidence helps classify the training data into clean and noisy subsets. Finally, DURA uses CEL and DSH loss functions to train the model, applying positive learning to the clean data and negative learning to the noisy data. In conclusion, the main contributions of this work are as follows:

- We propose Dynamic Uncertainty and Relational Alignment (DURA) framework to provide trusted retrieval in an effective and efficient way. Our DURA could be directly applied to robustly learn with noisy correspondence for Text-to-Image person search.

- A Dynamic Softmax Hinge Loss (DSH-Loss) is proposed to mitigate the adverse effects of unreliability caused by noisy correspondence. Specifically, DSH smoothly increases the difficulty of negative samples during training to improve the robustness against noisy correspondence.

- Extensive experiments verify that the proposed method improves the robustness against noisy correspondence, especially the high noise rate. Moreover, we provide insightful analysis that the learned uncertainty could reduce the negative impacts of noisy correspondences, improving the robustness.

2 Related Work

2.1 Text-to-Image Person Search

The concept of Text-to-Image Person Search was introduced by Li et al. [27] with the CUHK-PEDES dataset. Early methods used VGG and LSTM [27, 3] for feature extraction, aligning image-text features via matching loss functions. Later works [5, 15, 8] employed ResNet and BERT to enhance representations and design advanced matching losses. Current approaches are categorized into global-matching [34, 20, 42, 21], which align features in a unified space, and local-matching [6, 9, 22, 37], which focuses on fine-grained alignments like body parts and text elements for better retrieval.

Despite progress, most methods assume perfect training pairs, which is unrealistic due to noise. To address this, Yang et al. [38] proposed a robust learning method leveraging isovariant similarity consistency for noisy data. However, it depends on high-quality anchor points, and poor selection can lead to misclassification.

2.2 Uncertainty-based Learning

Many deep models[41] have been used with great success in applications, but they use deterministic predictions and lack the ability to assess their output uncertainty. To overcome this challenge, many researches [13, 18, 1] proposed methods to estimate the output uncertainty of the models. Kendall et al. [23] proposed the Bayesian deep learning framework combining input-related intrinsic uncertainty with knowledge-based uncertainty, focusing on the application of the model in per-pixel semantic segmentation and deep regression tasks. Chen et al. [4] simultaneously modelled the output uncertainty of the models by considering multi-granularity uncertainty for both coarse-grained and fine-grained image retrieval, integrating uncertainty modelling and uncertainty regularisation to improve recall and enhance the retrieval process. In contrast, our approach focuses more on modelling uncertainty across schema correspondences and aims to achieve robustness and reliability in text-to-image person search.

3 Methodology

In this section, we presents the proposed Dynamic Uncertainty and Relational Alignment (DURA) framework, which incorporates the Key Feature Selector (KFS) module and a novel loss function, the Dynamic Softmax Hinge Loss (DSH-Loss), to achieve robust cross-modal retrieval. The overview of DURA is illustrated in Fig. 1 and the details are discussed in the following subsections.

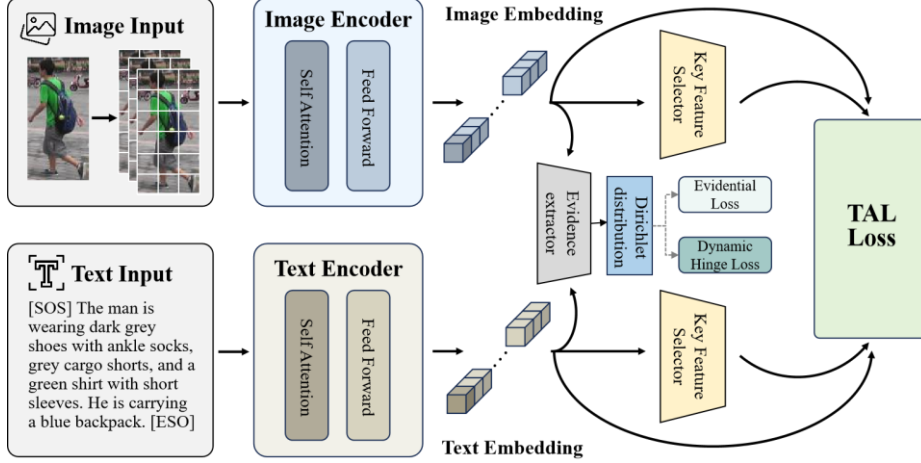


Fig. 1: The overview of our Dynamic Uncertainty and Relational Alignment framework (DURA), illustrating feature extraction, evidence-based uncertainty modeling, and evidence-guided training with TAL, CEL, DSH, and KFS.

3.1 Feature Extraction with Dual-Encoder

Our framework uses a dual-encoder architecture for feature extraction, utilizing pre-trained encoders for both image and text to ensure strong cross-modal alignment. Following the success of CLIP in aligning visual and textual embeddings, we initialize both the image and text encoders with the full CLIP architecture, improving the model’s ability to extract semantically aligned features from both modalities.

Formally, the gallery set is represented as $V = (I_i, y_i^p, y_i^v)_{i=1}^{N_v}$, and the corresponding textual description set is $T = (T_i, y_i^v)_{i=1}^{N_t}$, where N_v and N_t denote the number of images and texts, respectively. The label $y_i^p \in Y^p = 1, \dots, C$ is the class label (person identity), with C being the total number of identities, and $y_i^v \in Y^v = 1, \dots, N_v$ is the image label. The paired image-text dataset for TIPeID is defined as $P = (I_i, T_i, y_i^p, y_i^v)_{i=1}^N$, where each image-text pair shares the same image label y_i^v and class label y_i^p .

Image Encoder: Given an input image $I \in \mathbb{R}^{H \times W \times C}$, We use a CLIP pre-trained Vision Transformer (ViT) to extract image embeddings. An input image $I \in \mathbb{R}^{H \times W \times C}$ is divided into $N = \frac{H \times W}{p^2}$ non-overlapping patch size. Each patch is flattened and

transformed into a token representation $\{f_v^i\}_{i=1}^N$ with positional embeddings added and a special [CLS] token prepended. This tokenized sequence is processed through L-layer transformer blocks with self-attention, and the [CLS] token output f_t^{cls} captures the global image representation. A linear projection maps f_t^{cls} to a shared image-text embedding space for alignment with text embeddings.

Text Encoder: Given an input text T_i , We use a CLIP pre-trained transformer-based text encoder. Input text T_i is tokenized using byte pair encoding (BPE) into a sequence of subword tokens. Special tokens [SOS] and [EOS] are added to mark the boundaries, resulting in a sequence $\{f_t^{sos}, f_t^1, \dots, f_t^N, f_t^{eos}\}$. This sequence is processed through L-layer transformer blocks with self-attention, and the [EOS] token output f_t^{eos} captures the global text representation. A linear projection maps f_t^{eos} to the shared image-text embedding space, enabling alignment with image embeddings.

To measure the similarity between an image-text pair (I_i, T_j) , we utilize the global features represented by the [CLS] and [EOS] tokens. The global embedding similarity S_{ij} is computed using cosine similarity, defined as:

$$S_{ij} = \frac{f_v^{cls^T} \cdot f_t^{eos}}{\|f_v^{cls}\| \|f_t^{eos}\|} \quad (1)$$

Here, f_t^{eos} represents the global text embedding derived from the [EOS] token, while f_v^{cls} denotes the global image embedding obtained from the [CLS] token. These global embeddings are used to compute the similarity S_{ij} between the image and text in the shared multimodal embedding space. Additionally, $\{f_t^j\}_{j=1}^N$ and $\{f_v^j\}_{j=1}^N$ correspond to the local features extracted from the N word token in the text T_i and the N patches in the image I_i , respectively, providing fine-grained contextual representations.

3.2 Key Feature Selector

While global embeddings obtained from [CLS] and [EOS] tokens capture high-level similarities between image-text pairs, they often overlook subtle, fine-grained details essential for accurate text-to-image person search. To address this issue, we introduce a Key Feature Selector (KFS) module that enhances the discriminative power of the learned representations by incorporating informative local features.

We begin by applying L2 normalization to both visual and textual features: $f_v^i = \text{L2Norm}(f_v^i)$ and $f_t^i = \text{L2Norm}(f_t^i)$. This normalization step ensures consistent magnitudes across features, improving their stability and reliability, especially under noisy conditions.

Next, we refine features using Max-K pooling, which selects the top k values and averages them. This process emphasizes the most discriminative components, allowing the model to focus on the critical cues that distinguish one identity from another.

Before pooling, we further enhance feature representations using a combination of MLP, FC, and a Squeeze-and-Excitation (SE) layer at the MLP's output. The MLP

and FC layers transform the input features into a richer, more expressive space, while the SE layer adaptively recalibrates channel-wise feature responses, highlighting informative dimensions and suppressing less relevant ones.

Formally, the visual and textual features F_v^i and F_t^i are computed as:

$$F_v^i = \text{MaxPool}(\text{SE}(\text{MLP}(V_i^s)) + \text{FC}(f_v^i)), \quad (2)$$

$$F_t^i = \text{MaxPool}(\text{SE}(\text{MLP}(T_i^s)) + \text{FC}(f_t^i)) \quad (3)$$

By integrating global and local perspectives within DURA, the KFS module ensures a balanced representation that captures both high-level alignment and fine-grained distinctions. This leads to improved robustness and accuracy in text-to-image person search, particularly in complex and noisy retrieval scenarios.

3.3 Uncertainty Modeling

In this section, we model cross-modal uncertainty using the Dempster-Shafer Theory of Evidence, guided by the principles of Subjective Logic. Consider a mini-batch of N image-text pairs. For each pair (I_i, T_j) , we first compute a similarity score S_{ij} . To translate this score into evidence, we apply an evidence extractor $f(\cdot)$:

$$e_{ij} = f(S_{ij}) = \exp(\tanh(S_{ij}/T)), \quad (4)$$

where $0 < T < 1$ is a scaling parameter. Thus, for a given visual query I_i , the evidence vector $e_i \rightarrow T_i$ can be extracted from the cross-modal similarities using Equation (4), i.e., $e_i \rightarrow T_i = [e_{i1}, e_{i2}, \dots, e_{iN}]$. Similarly, the evidence vector $e_{T_2 \rightarrow i}$ for a given textual query T_i can be obtained as: $e_{T_2 \rightarrow i} = [e_{1i}, e_{2i}, \dots, e_{Ni}]$.

Subjective Logic assigns a belief mass to each query and an overall uncertainty mass based on the collected cross-modal evidence (e.g., $e_{i \rightarrow T_i}$ and $e_{T_i \rightarrow i}$), which can be defined as:

$$b_{ij} = \frac{e_{ij}}{L_i} \text{ and } u_i = \frac{N}{L_i}, \quad (5)$$

where: $L_i = \sum_{j=1}^N (e_{ij} + 1)$ and $u_i + \sum_{j=1}^N b_{ij} = 1$. The term L_i represents the strength of the Dirichlet distribution, while the belief mass assignment $b_i = [b_{i1}, b_{i2}, \dots, b_{iN}]$ corresponds to subjective opinions derived from the Dirichlet distribution with parameters $\alpha_i = [\alpha_{i1}, \alpha_{i2}, \dots, \alpha_{iN}]$, where: $\alpha_{ij} = e_{ij} + 1$.

Intuitively, cross-modal retrieval is analogous to classifying instances (i.e., pairs), where the query similarity aligns with the probability assignment. The Dirichlet distribution, parametrized by the evidence, represents the density of these probability assignments. Thus, α_i models second-order probabilities and uncertainties. The density function is defined as:

$$D(p_i | \alpha_i) = \begin{cases} \frac{1}{B(\alpha_i)} \prod_{j=1}^N p_{ij}^{\alpha_{ij}-1}, & \text{for } p_i \in \mathcal{S}_N, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $p_i \in \mathbb{R}^N$ are the query probabilities, $B(\alpha_i)$ is the N-dimensional multinomial beta function, and S_N denotes the N-dimensional unit simplex .

In this manner, the evidence extracted from cross-modal similarities informs a second-order distribution (the Dirichlet), representing both the likelihood of each potential match and the associated uncertainty. This approach ensures a more robust handling of noisy and uncertain retrieval scenarios.

3.4 Cross-modal Evidential Learning

Following the approach of [33], we treat cross-modal retrieval as a K-way classification task, where each query is expected to match its corresponding target. Let y_i be a K-dimensional one-hot vector indicating the correct match for query i . The model produces evidence $\alpha_i = [\alpha_{i1}, \dots, \alpha_{iK}]$ that parameterizes a Dirichlet distribution, modeling both the probability assignments and their associated uncertainty.

To align the model's probability estimates $\mathbb{E}[p_{ij}] = \frac{\alpha_{ij}}{L_i}$ (with $L_i = \sum_{k=1}^K \alpha_{ik}$) to the ground truth y_i , we employ the following mean-squared loss over the Dirichlet distribution:

$$\mathcal{L}_m(\alpha_i, y_i) = \sum_{j=1}^K [(y_{ij} - \alpha_{ij})^2 + \frac{\alpha_{ij}(L_i - \alpha_{ij})}{L_i^2(L_i + 1)}] \quad (7)$$

Minimizing \mathcal{L}_m encourages the expected probabilities to approximate the ground truth while reducing uncertainty.

However, \mathcal{L}_m alone does not ensure that evidence for negative matches diminishes. To address this, we introduce a Kullback-Leibler (KL) divergence term that penalizes excessive evidence for incorrect targets. We define:

$$\tilde{\alpha}_i = y_i + (1 - y_i) \odot \alpha_i, \quad (8)$$

and measure the divergence from a uniform Dirichlet distribution $D(p_i|1)$:

$$\mathcal{L}_{KL}(\alpha_i, y_i) = \text{KL}(D(p_i|\tilde{\alpha}_i) \parallel D(p_i|1)). \quad (9)$$

This term encourages the model to reduce unwarranted evidence for mismatched pairs, enhancing robustness against noisy matches.

For an image query I_i , the evidential loss is:

$$\mathcal{L}_e^{i2t}(I_i, l_i) = \mathcal{L}_m(\alpha_i^{i2t}, l_i) + \lambda_2 \mathcal{L}_{KL}(\alpha_i^{i2t}, l_i), \quad (10)$$

where λ_2 is a balance factor. Similarly, we define $\mathcal{L}_e^{2i}(T_i, l_i)$ for text-to-image retrieval. Combining them yields a bidirectional evidential loss:

$$\mathcal{L}_e(I_i, T_i, l_i) = \mathcal{L}_e^{i2t}(I_i, l_i) + \mathcal{L}_e^{2i}(T_i, l_i). \quad (11)$$

In essence, this evidential learning framework provides a principled way to handle uncertain and noisy matches in cross-modal retrieval. By treating retrieval as a classification problem under a Dirichlet prior, we can jointly optimize

for accuracy and robustness, ensuring that the model not only fits the correct matches but also avoids overconfidence in incorrect ones.

3.5 Dynamic Softmax Hinge Loss

Conventional hinge-based losses [11] often consider all negative samples in a mini-batch, which can amplify errors when noisy correspondences are present. This overemphasis on all negatives may reduce model robustness and lead to unstable training. To mitigate this issue, Qin et al. [33] proposed the Robust Dynamic Hinge (RDH) loss, focusing only on the single hardest negative. While this approach improves stability under noise, it overlooks other potentially informative negatives, limiting the model’s ability to learn from the broader negative distribution.

To achieve a more balanced and stable training process, we propose the Dynamic Softmax Hinge (DSH) loss, which adaptively utilizes a controlled subset of hard negatives rather than focusing solely on the single hardest one. By doing so, DSH maintains robustness against noise while leveraging a richer set of negative examples. This encourages a more comprehensive and effective representation learning process.

$$\begin{aligned} \mathcal{L}_h(I, T) = & \frac{1}{n} \left[\sum_{j=1}^n \left[\gamma - S(I, T) + \tau \cdot \log \left(\sum_{j=1}^N \exp \left(\frac{S(I, T_j)}{\tau} \right) \right) \right] \right]_+ \\ & + \sum_{j=1}^n \left[\gamma - S(I, T) + T \cdot \log \left(\sum_{j=1}^N \exp \left(\frac{S(I_j, T)}{\tau} \right) \right) \right]_+ \end{aligned} \quad (12)$$

where $[x]_+ = \max(x, 0)$, n is the dynamically adjusted number of hardest negatives, and γ is the margin. The value of n decreases dynamically during training:

$$n = \max(\lfloor K - \eta \cdot \text{Step} \rfloor, \mu), \quad (13)$$

where K is the mini-batch size, η is the annealing coefficient, Step is the current training step, and μ is the lower bound of n .

3.6 Cross-Modal Alignment

Triplet Ranking Loss (TRL), introduced by Schroff et al. in 2015, has been widely applied in cross-modal alignment tasks, leveraging the hardest negative samples to distinguish positive and negative pairs effectively. However, TRL suffers from performance limitations due to its insufficient focus on the broader distribution of negative samples, which can lead to suboptimal convergence and instability. To address this, Qin et al. [32] proposed the Triplet Alignment Loss (TAL) in 2024, which relaxes the optimization from focusing solely on the hardest negative to considering all negative samples with an upper bound constraint. This relaxation reduces the risk of optimization being dominated by the hardest negatives, enhancing the stability of the training process while maintaining a comprehensive utilization of negative samples. By balancing attention to both hard and overall negative samples, TAL achieves more stable and robust training, ensuring better performance in cross-modal alignment tasks. Therefore, this paper adopts TAL as the loss function for cross-modal alignment.

For an input pair (I_i, T_i) in a mini-batch x , TAL is defined as:

$$\begin{aligned} \mathcal{L}_{\text{TAL}}(I_i, T_i) = & [m - S_{i \rightarrow t}^+(I_i) + \tau \log(\sum_{j=1}^K \exp(\frac{S(I_i, T_j)}{\tau}))]^+ \\ & + [m - S_{t \rightarrow i}^+(T_i) + \tau \log(\sum_{j=1}^K \exp(\frac{S(I_j, T_i)}{\tau}))]^+ \end{aligned} \quad (14)$$

where m is a positive margin coefficient, τ is a temperature coefficient to control hardness, $S(I_i, T_j) \in \{S_{ij}^b, S_{ij}^t\}$, $[x]^+ \equiv \max(x, 0)$, $\exp(x) \equiv e^x$, and K is the size of x . From Lemma 1, as $T \rightarrow 0$, TAL approaches TRL and focuses more on hard negatives. Since multiple positive pairs from the same identity may appear in the mini-batch, is the weighted average similarity of positive pairs for image I_i

$$S_{i \rightarrow t}^+(I_i) = \sum_{j=1}^K \alpha_{ij} S(I_i, T_j) \quad (15)$$

3.7 Overall Loss Function

DURA aims to enhance global image-text representations in a joint embedding space by integrating multiple objectives. Specifically, it utilizes an evidential loss (\mathcal{L}_e) for handling noisy correspondences, a Dynamic Softmax Hinge loss (\mathcal{L}_h) for controlling the difficulty of negative samples, and a Triplet Alignment Loss (\mathcal{L}_{TAL}) for stable and comprehensive cross-modal alignment. Together, these losses foster fine-grained feature interaction, uncertainty modeling, and robust image-text matching.

We train DURA in an end-to-end fashion, combining all components into a single optimization objective:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_e(I, T, l) + \mathcal{L}_h(I, T) + \mathcal{L}_{\text{TAL}}(I, T) \quad (10)$$

4 Experiments

In this section, we evaluate our DURA framework on three widely used text-to-image person search datasets—CUHK-PEDES, ICFG-PEDES—under various noise conditions to demonstrate its robustness.

4.1 Datasets and Performance Measurements

CUHK-PEDES [28], the first text-to-image person retrieval dataset, includes 40,206 images and 80,412 descriptions for 13,003 identities. Split: training (11,003 identities, 34,054 images, 68,108 descriptions), validation (1,000 identities, 3,078 images, 6,158 descriptions), test (1,000 identities, 3,074 images, 6,156 descriptions). Descriptions average 23 words.

ICFG-PEDES [9] has 54,522 images for 4,102 identities, each with one 37-word description. Training: 3,102 identities, 34,674 image-text pairs. Test: 1,000 identities, 19,848 pairs. Notable for one-to-one image-text pairing.

Evaluation Metrics: Rank-k ($k=1, 5, 10$) measures top-k retrieval accuracy. Mean Average Precision (mAP) and mean Inverse Negative Penalty (mINP) [39] assess overall performance. Higher values indicate better results.

4.2 Implementation Details

DURA uses CLIP-ViT-B/16 as the image encoder and CLIP text Transformer as the text encoder, with a multimodal interaction encoder (512 hidden size, 8 attention heads). Data augmentations include random horizontal flip, crop, and erasing; a Key Feature Selector uses a 0.5 sampling ratio. Images are resized to 384×128 , and text sequences are truncated to 77 tokens.

Training spans 60 epochs with the Adam optimizer, a starting learning rate of 8×10^{-6} , cosine decay, and a 2-epoch warm-up. Following RDE, TAL’s margin is 0.1, and temperature τ is 0.015. Experiments run on a single RTX 2080 Ti GPU, ensuring robust evaluation of DURA’s noise resistance and cross-modal feature extraction.

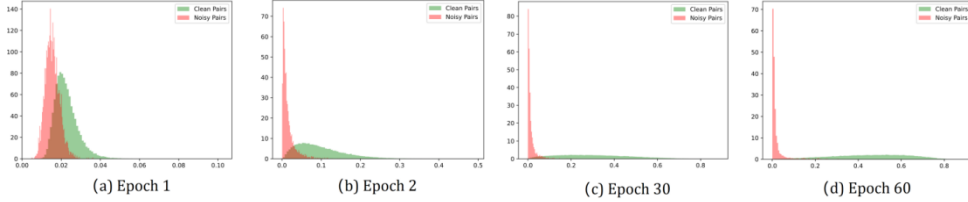


Fig. 2: We visualize the evidence distribution of clean and noisy pairs at different training stages of our DURA, which is conducted on CHUKPEDES under 20% noise.

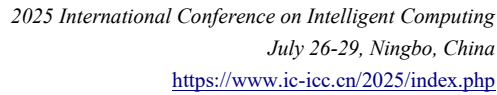
4.3 Comparison with Noise Correspondence

We compare our DURA model with six state-of-the-art baselines: SSAN [9], IVT [34], IRRa [19], DECL [33], RDE [32], and CLIP-C, using CLIP-ViT-B/16 as the baseline. Experiments evaluate text-to-image person search on CUHK-PEDES, ICFG-PEDES, and RSTPReid datasets, with noisy correspondences (0%, 20%, 50%) introduced by shuffling images. For rigor, RDE was rerun on an RTX 2080 Ti 12GB GPU, denoted RDE*.

CUHK-PEDES: Table 1 shows DURA’s superior robustness to noise. At 20% noise, DURA achieves 75.04% (Rank-1), 89.74% (Rank-5), 93.66% (Rank-10), 66.81% (mAP), and 50.48% (mINP), outperforming all baselines. DURA maintains strong performance across low and high noise, with minimal degradation, highlighting its resilience.

ICFG-PEDES: DURA’s dominance on the 1K test set. At 50% noise, DURA improves Rank-1 by up to 12% over baselines, excelling in all metrics (Rank-1, Rank-5, Rank-10, mAP, mINP) under both low and high noise, demonstrating robust retrieval.

RSTPReid: DURA’s exceptional performance at high noise, achieving 62.94% (Rank-1), 83.55% (Rank-5), 89.45% (Rank-10), 47.92% (mAP), and 25.55% (mINP). RDE* ranks second, but DURA consistently leads, showcasing adaptability to noisy conditions.



Noise	Methods	CUHK-PEDES						ICFG-PEDES				
		Type	R-1	R-5	R-10	mAP	mINP	R-1	R-5	R-10	mAP	mINP
0%	SSAN	Best	61.37	80.15	86.73	-	-	54.23	72.63	79.53	-	-
	IVT	Best	65.59	83.11	89.21	-	-	56.04	73.60	80.22	-	-
	CFine	Best	69.57	85.93	91.15	-	-	60.83	76.55	82.42	-	-
	IRRA	Best	73.38	89.93	93.71	66.13	50.24	63.46	80.25	85.82	38.06	7.93
	DECL	Best	71.36	88.11	92.48	64.25	48.26	63.42	79.29	84.89	37.02	6.57
	RDE [*]	Best	75.76	90.14	94.18	67.56	51.43	67.68	82.47	87.36	40.06	7.87
	DURA	Best	76.14	90.42	94.25	67.68	51.57	67.88	82.67	87.66	40.37	8.21
20%	SSAN	Best	46.52	68.36	77.42	42.49	28.13	40.57	62.58	71.53	20.93	2.22
		Last	45.76	67.98	76.28	40.05	24.12	40.28	62.68	71.53	20.98	2.25
	IVT	Best	58.59	78.51	85.61	57.19	45.78	50.21	69.14	76.18	34.72	8.77
		Last	57.67	78.04	85.02	56.17	44.42	48.70	67.42	75.06	34.44	9.25
	IRRA	Best	69.74	87.09	92.20	62.28	45.84	60.76	78.26	84.01	35.87	6.80
		Last	69.44	87.09	92.04	62.16	45.70	60.58	78.14	84.20	35.92	6.91
	CLIP-C	Best	66.41	85.15	90.89	59.36	43.02	55.25	74.76	81.32	31.09	4.94
		Last	66.10	86.01	91.02	59.77	43.57	55.17	74.58	81.46	31.12	4.97
	DECL	Best	70.29	87.04	91.93	62.84	46.54	61.95	78.36	83.88	36.08	6.25
		Last	70.08	87.20	92.14	62.86	46.63	61.95	78.36	83.88	36.08	6.25
	RDE [*]	Best	74.46	89.10	<u>93.63</u>	66.13	49.66	66.54	81.70	86.70	39.08	7.55
		Last	74.53	89.23	93.55	66.13	49.63	66.51	81.70	86.71	39.09	7.56
	DURA	Best	75.04	89.74	93.66	66.81	50.48	66.62	81.96	86.86	39.53	7.72
		Last	<u>74.60</u>	<u>89.44</u>	<u>93.45</u>	<u>66.43</u>	<u>50.23</u>	<u>66.53</u>	<u>81.72</u>	<u>86.72</u>	<u>39.34</u>	<u>7.68</u>
50%	SSAN	Best	13.43	31.74	41.89	14.12	6.91	18.83	37.70	47.43	9.83	1.01
		Last	11.31	28.07	37.90	10.57	3.46	17.06	37.18	47.85	6.58	0.39
	IVT	Best	50.49	71.82	79.81	48.85	36.60	43.03	61.48	69.56	28.86	6.11
		Last	42.02	65.04	73.72	40.49	27.89	36.57	54.83	62.91	24.30	5.08
	IRRA	Best	62.41	82.23	88.40	55.52	38.48	52.53	71.99	79.41	29.05	4.43
		Last	42.79	64.31	72.58	36.76	21.11	39.22	60.52	69.26	19.44	1.98
	CLIP-C	Best	64.02	83.66	89.38	57.33	40.90	51.60	71.89	79.31	28.76	4.33
		Last	63.97	83.74	89.54	57.35	40.88	51.49	71.99	79.32	28.77	4.37
	DECL	Best	65.22	83.72								

Table 4: Ablation studies on the CUHK-PEDES dataset

No.	Methods	R-1	R-5	R-10	mAP	mINP
0	Baseline	66.41	85.15	90.89	59.36	43.02
1	+TAL	70.45	86.74	91.52	62.33	45.77
2	+TAL+L _h	71.05	87.15	91.80	63.26	46.89
3	+TAL+L _e	70.74	87.22	91.85	63.25	46.85
4	+TAL+KFS	73.44	88.87	92.32	65.24	48.69
5	+TAL+KFS+L _e	74.58	88.97	92.73	66.09	49.81
6	+TAL+KFS+L _h	74.08	88.74	92.52	65.96	49.09
7	+TAL+KFS+L _e +L _h	75.04	89.74	93.66	66.81	50.48

4.4 Ablation Study

We performed an ablation study on CUHK-PEDES with 20% noise, using CLIP-ViT-B/16 as the baseline to evaluate DURA’s components. Table 4 shows: (1) TAL alone (No.0 vs. No.1) boosts Rank-1 (+4.04%), Rank-5 (+1.59%), Rank-10 (+0.63%), mAP (+2.97%), and mINP (+2.75%). (2) Adding KFS atop TAL (No.1 vs. No.4, No.2 vs. No.6, No.3 vs. No.5) significantly improves all metrics. (3) Le or Lh with TAL and KFS (No.4 vs. No.5, No.4 vs. No.6) partially enhances performance. (4) Full DURA (No.7) achieves top results: 75.04% (Rank-1), 89.74% (Rank-5), 93.66% (Rank-10), 66.81% (mAP), and 50.48% (mINP), confirming robust component synergy.

5 Conclusion

In this work, we present the Dynamic Uncertainty and Relational Alignment (DURA) framework, aimed at improving text-based person search under noisy conditions. DURA leverages a Key Feature Selector (KFS) and a Dynamic Soft-max Hinge Loss (DSH-Loss) to robustly handle noisy training pairs, while its cross-modal evidential learning distinguishes clean from noisy data. Experiments on CUHK-PEDES, ICFG-PEDES, and RSTPReid confirm that DURA consistently outperforms existing methods, especially in high-noise scenarios, demonstrating its practical effectiveness when data quality is limited.

References

1. Soyed Tuhin Ahmed. Scalable and efficient methods for uncertainty estimation and reduction in deep learning, 2024.
2. Yang Bai, Min Cao, Daming Gao, Ziqiang Cao, Chen Chen, Zhenfeng Fan, Liqiang Nie, and Min Zhang. Rasa: Relation and sensitivity aware representation learning for text-based person search. In Proceedings of the Thirty-Second International Joint Conference



- on Artificial Intelligence, IJCAI-2023, page 555–563. International Joint Conferences on Artificial Intelligence Organization, August 2023.
3. Tianlang Chen, Chenliang Xu, and Jiebo Luo. Improving text-based person search by spatial matching and adaptive threshold. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), pages 1879–1887, 2018.
 4. Yiyang Chen, Zhedong Zheng, Wei Ji, Leigang Qu, and Tat-Seng Chua. Composited image retrieval with text feedback via multi-grained uncertainty regularization, 2024.
 5. Yucheng Chen, Rui Huang, Hong Chang, Chuanqi Tan, Tao Xue, and Bingpeng Ma. Cross-modal knowledge adaptation for language-based person search. IEEE Transactions on Image Processing, page 4057–4069, Jan 2021.
 6. Yuhao Chen, Guoqin Zhang, Yujiang Lu, Zhnxing Wang, Yuhui Zheng, and Ruili Wang. Tipcb: A simple but effective part-based convolutional baseline for text-based person search, 2021.
 7. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018.
 8. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North, Jan 2019.
 9. Zefeng Ding, Changxing Ding, Zhiyin Shao, and Dacheng Tao. Semantically self-aligned network for text-to-image part-aware person re-identification, 2021.
 10. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, and et al. Sylvain Gelly. An image is worth 16x16 words: Transformers for image recognition at scale, 2020.
 11. Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. arXiv preprint arXiv:1707.05612, 2017.
 12. Yanglin Feng, Hongyuan Zhu, Dezhong Peng, Xi Peng, and Peng Hu. Rono: Robust discriminative learning with noisy labels for 2d-3d cross-modal retrieval. In 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11610–11619, 2023.
 13. Gianni Franchi, Olivier Laurent, Maxence Legu  ry, Andrei Bursuc, Andrea Pilzer, and Angela Yao. Make me a bnn: A simple strategy for estimating bayesian uncertainty from pre-trained models, 2023.
 14. Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. Neural Information Processing Systems, Neural Information Processing Systems, Jan 2018.
 15. Xiao Han, Sen He, Li Zhang, and Tao Xiang. Text-based person search with limited data, 2021.
 16. Peng Hu, Xi Peng, Hongyuan Zhu, Liangli Zhen, and Jie Lin. Learning cross-modal retrieval with noisy labels. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5399–5409, 2021.
 17. Yan Huang, Zhang Zhang, Qiang Wu, Yi Zhong, and Liang Wang. Attribute-guided pedestrian retrieval: Bridging person re-id with internal attribute variability. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 17689–17699, 2024.

18. Joel Jaskari, Jaakko Sahlsten, Theodoros Damoulas, Jeremias Knoblauch, Simo Särkkä, Leo Kärkkäinen, Kustaa Hietala, and Kimmo Kaski. Uncertainty-aware deep learning methods for robust diabetic retinopathy classification, 2022.
19. D. Jiang and M. Ye. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2787–2797, 2023.
20. Ding Jiang and Mang Ye. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval, 2023.
21. Ding Jiang and Mang Ye. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2787–2797, 2023.
22. Ya Jing, Chenyang Si, Junbo Wang, Wei Wang, Liang Wang, and Tieniu Tan. Pose-guided multi-granularity attention network for text-based person search, 2019.
23. Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision?, 2017.
24. Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
25. Jie Lei, Xinlei Chen, Ning Zhang, Mengjiao Wang, Mohit Bansal, Tamara L. Berg, and Licheng Yu. Loopitr: Combining dual and cross encoder architectures for image-text retrieval, 2022.
26. Shenshen Li, Xing Xu, Yang Yang, Fumin Shen, Yijun Mo, Yujie Li, and Heng Tao Shen. Dcel: Deep cross-modal evidential learning for text-based person retrieval. *Proceedings of the 31st ACM International Conference on Multimedia*, 2023.
27. Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person search with natural language description. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5187–5196, 2017.
28. Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. Person search with natural language description. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5187–5196, 2017.
29. Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1487–1495, 2019.
30. Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, SarahM. Erfani, and JamesE. Bailey. Normalized loss functions for deep learning with noisy labels. *arXiv: Learning*, arXiv: Learning, Jun 2020.
31. Antoine Miech, Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. Thinking fast and slow: Efficient text-to-visual retrieval with transformers. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9821–9831, 2021.
32. Yaqing Qin, Yu Chen, Donghao Peng, et al. Noisy-correspondence learning for text-to-image person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27197–27206, 2024.
33. Yaqing Qin, Donghao Peng, Xi Peng, et al. Deep evidential learning with noisy correspondence for cross-modal retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4948–4956. ACM, 2022.
34. Xiujun Shu, Wei Wen, Haoqian Wu, Keyu Chen, Yiran Song, Ruizhi Qiao, Bo Ren, and Xiao Wang. See finer, see more: Implicit modality alignment for text-based person retrieval, 2022.



35. Siqi Sun, Yen-Chun Chen, Linjie Li, Shuohang Wang, Yuwei Fang, and Jingjing Liu. Lightningdot: Pre-training visual-semantic embeddings for real-time image-text retrieval, 2021.
36. Haochen Wang, Jiayi Shen, Yongtuo Liu, Yan Gao, and Efstratios Gavves. Nformer: Robust person re-identification with neighbor transformer. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 7287–7297, 2022.
37. Zhe Wang, Zhiyuan Fang, Jun Wang, and Yezhou Yang. Vitaa: Visual-textual attributes alignment in person search by natural language, 2020.
38. Yuchen Yang, Likai Wang, Erkun Yang, and Cheng Deng. Robust noisy correspondence learning with equivariant similarity consistency. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 17700–17709, 2024.
39. Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi. Deep learning for person re-identification: A survey and outlook. IEEE Transactions on Pattern Analysis and Machine Intelligence, page 2872–2893, Jun 2022.
40. Huaiwen Zhang, Yang Yang, Fan Qi, Shengsheng Qian, and Changsheng Xu. Robust video-text retrieval via noisy pair calibration. IEEE Transactions on Multimedia, 25:8632–8645, 2023.
41. Ji Zhang, Jingkuan Song, Lianli Gao, Ye Liu, and Heng Tao Shen. Progressive meta-learning with curriculum. IEEE Transactions on Circuits and Systems for Video Technology, 32(9):5916–5930, Sep 2022.
42. Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, Computer Vision – ECCV 2018, pages 707–723, Cham, 2018. Springer International Publishing.
43. Aichun Zhu, Zijie Wang, Yifeng Li, Xili Wan, Jing Jin, Tian Wang, Fangqiang Hu, and Gang Hua. Dssl: Deep surroundings-person separation learning for text-based person retrieval. In Proceedings of the 29th ACM International Conference on Multimedia, Oct 2021.