

HIERARCHICAL ENCODING TREE WITH MODALITY MIXUP FOR CROSS-MODAL HASHING

Zhiping Xiao¹, Junyu Luo^{2†}, Hang Zhou³, Yusheng Zhao², Xiao Luo^{4†}, Pengyun Wang⁵, Wei Ju², Siyu Heng⁶, Ming Zhang^{2†}

¹ Paul G. Allen School of Computer Science and Engineering, University of Washington

² State Key Laboratory for Multimedia Information Processing, School of Computer Science, PKU-Anker LLM Lab, Peking University

³ Department of Statistics and Operations Research & School of Data Science and Society, University of North Carolina at Chapel Hill

⁴ Department of Statistics, University of Wisconsin–Madison

⁵ Data Science Institute, University of Chicago

⁶ Department of Biostatistics, New York University
patxiao@uw.edu, luo.junyu@outlook.com,
xiao.luo@wisc.edu, mzhang_cs@pku.edu.cn

ABSTRACT

Cross-modal retrieval is a fundamental task that aims to learn semantic correspondences across different data modalities, such as visual and textual modalities. Unsupervised hashing methods can efficiently manage large-scale data and can be effectively applied to cross-modal retrieval studies. However, existing methods typically fail to fully exploit the hierarchical semantic structure within text and image data, where instances naturally organize into multi-level communities of varying granularity. Moreover, the commonly-used direct modal alignment cannot effectively bridge the semantic gap between these two modalities. To address these issues, we introduce a novel **H**ierarchical **E**ncoding **T**ree with **M**odality **M**ixup (**HINT**) method, which achieves effective cross-modal retrieval by extracting hierarchical cross-modal relations. HINT constructs a cross-modal encoding tree guided by hierarchical structural entropy and generates *proxy* samples of text and image modalities for each instance from the encoding tree. Through the curriculum-based mixup of proxy samples, HINT achieves progressive modal alignment and effective cross-modal retrieval. We also conduct cross-modal consistency learning to achieve global-view semantic alignment between text and image representations. Extensive experiments on a range of cross-modal retrieval datasets demonstrate the superiority of HINT over state-of-the-art methods.

1 INTRODUCTION

Cross-modal retrieval aims to measure the semantic similarity between different modalities, using retrieval methods such as approximate nearest neighbors (ANNs) search (Zhu et al., 2023; Zhen et al., 2019). Cross-modal retrieval has significant application value, such as in retrieval-augmented generation (RAG) (Li et al., 2024b; Cui et al., 2024) and search engines (Song et al., 2024; Chen et al., 2017). With the rapid development of large-scale vision-language datasets, cross-modal retrieval has attracted increasing attention. Therefore, researchers have turned to hashing-based cross-modal retrieval (Cao et al., 2017; Yan et al., 2020), which achieves efficient storage and indexing by substituting bit-wise operations for computationally prohibitive pairwise distance comparisons. Hashing-based cross-modal retrieval methods map high-dimensional semantic vectors from different modalities into a unified Hamming space (binary space) (Luo et al., 2023; Huang et al., 2024), enabling similarity comparison and ANNs.

[†]Corresponding authors.

Cross-modal hashing has garnered significant attention from the community (Zhang et al., 2024b; Sun et al., 2024; Luo et al., 2025). It includes supervised and unsupervised methods. Supervised cross-modal hashing (Shen et al., 2024; Ma et al., 2024; Liu et al., 2019a; Lu et al., 2019) learn hash codes using labeled data. However, due to the expensiveness and scarcity of cross-modal annotations in real-world scenarios (Wang et al., 2023), researchers have shifted their focus to unsupervised methods that do not rely on annotations. Unsupervised cross-modal hashing (Liang et al., 2024; Li et al., 2024a; Zhang et al., 2018) leverage pair-wise cross-modal data, exploiting the similarity between samples and employing mechanisms such as contrastive learning (Hu et al., 2022) and adversarial learning (Li et al., 2019) to guide hash learning.

Unsupervised cross-modal hashing has made promising progress (Liang et al., 2024; Li et al., 2024a; Zhang et al., 2024b; Tu et al., 2023), but still suffers from the following issues: The first challenge is *the lack of hierarchical semantic structure*. Without semantic annotations (e.g., category labels), previous unsupervised methods can only leverage paired data (Hu et al., 2022) (e.g., image-text pairs) as supervision, which provides flat (i.e., no hierarchical structure, all data points are on the same level) and sparse signals. However, real-world text and image data naturally exhibit a hierarchical semantic structure, containing numerous local communities. The instances within each community have similar semantics, while the semantic differences across communities are substantial. The absence of hierarchical-information-mining leads to insufficient exploration of community relationships, hindering the learning of generalizable hash codes. Furthermore, this challenge is compounded by *the ineffective alignment of heterogeneous modalities*. Existing methods (Hu et al., 2022; Zhang et al., 2024b; Tu et al., 2023) employ different encoders to project data from various modalities and optimize towards a common objective. Nevertheless, due to the inherent heterogeneity across modalities (e.g., manifested in structure and semantics), direct alignment can pose a high learning difficulty and lead to suboptimal performance. Therefore, it is necessary to conduct hierarchical cross-modal learning and in a progressive manner.

To address the aforementioned issues, we propose a novel approach called **Hierarchical Coding Tree with Modality Mixup (HINT)** for hashing-based cross-modal retrieval. The core idea of HINT lies in constructing a cross-modal encoding tree that recovers hierarchical semantic structures and mines local semantic communities. Specifically, guided by the hierarchical structural entropy (Li et al., 2018; Zou et al., 2023), we construct the encoding tree from the enhanced cross-modal relationship graph.

The encoding tree has dense connections within local communities and sparse connections between communities. By utilizing the cross-modal encoding tree, we mitigate the performance degradation associated with flat and sparse cross-modal connections. Next, we synthesize proxy samples in different modalities for each sample based on the encoding tree. Through curriculum-based mixup on these proxy samples, we achieve progressive modality alignment, circumventing the challenging task of directly aligning heterogeneous modalities. We also achieve semantic alignment from a global perspective by optimizing the consistency of the semantic distributions of the proxy samples in different modalities. Extensive experiments on benchmark datasets demonstrate the superior performance of HINT.

The main contributions of this paper are: **1 Hierarchical Modeling Approach.** We connect the encoding tree with cross-modal hashing problems. Specifically, we construct a cross-modal encoding tree to explore the cross-modal relationships and uncover local semantic communities hierarchically. **2 Novel Integrated Framework.** Based on the hierarchical encoding tree, we extract cross-modal proxy samples. Leveraging the proxy samples, we design a curriculum-based modality mixup mechanism for effective cross-modal hash learning. These components synergistically form an end-to-end framework where each enables and enhances the others. On the other hand, we achieve global-view consistency learning through the distribution alignment. **3 Competitive Performance.** Comprehensive experiments demonstrate that HINT achieves state-of-the-art performance.

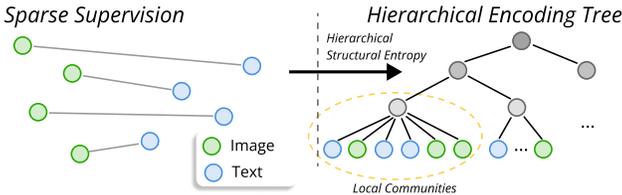


Figure 1: HINT transforms sparse cross-modal supervision (left) into a meaningful hierarchical encoding tree (right), which reveals local semantic communities for robust cross-modal alignment.

2 RELATED WORKS

Cross-modal Retrieval is a fundamental task for bridging data of different modalities (Lee et al., 2024; Chen et al., 2024; Li et al., 2023; Krojer et al., 2022; Radford et al., 2021; Ge et al., 2024). Due to the diverse distributions and structures of texts and images, it is necessary to map them effectively into a unified representation space to calculate the semantic similarity between samples (Ding et al., 2016b; Liu et al., 2019b). With this unified representation, we can employ the approximate nearest neighbors (ANNs) (Zhu et al., 2023; Zhen et al., 2019) methods for similarity search. Recently, researchers turn to cross-modal hashing methods to enhance efficiency for both storage costs and computational costly retrieval processes (Xu et al., 2017; Jiang & Li, 2019; Wang et al., 2024b).

Unsupervised Cross-modal Hashing (Liang et al., 2024; Li et al., 2024a; Zhang et al., 2024b; Wang et al., 2024b) utilizes data correlation information to project cross-modal data into a unified binary space (Huang et al., 2024; Wu et al., 2018). Due to the expensive and difficult acquisition of labeled cross-modal data, supervised methods sometimes face challenges in the real-world (Hu et al., 2022). Therefore, unsupervised cross-modal hashing has attracted widespread attention (Zhou et al., 2014; Gao et al., 2023; Mikriukov et al., 2022). Researchers also use adversarial networks (Li et al., 2019) and contrastive learning (Hu et al., 2022) to handle cross-modal hash learning. These methods generally rely on sparse text-image relationships, lacking local community mining. We explore the hierarchical cross-modal relationship and learn more generalizable hash representations through modality mixup and cross-modal consistency learning.

Cross-modal Relationship Modeling is an essential topic in multi-modal research (Oh et al., 2023; Wang et al., 2024a; Li et al., 2024c; Liang et al., 2022; Huang et al., 2021). Some methods employ similarity modeling within modalities (Zhang et al., 2018), such as constructing graph structures for images and texts separately using the Wasserstein metric. The tree-based methods (Ge et al., 2021; Chen et al., 2022) are also introduced for cross-modal relationship modeling. Recent hierarchical methods (Jin et al., 2021; Zhang et al., 2025) construct predefined multi-level structures using external tools or preset granularities for supervised retrieval tasks. Other works (Sun et al., 2023) adopt architectural hierarchies through progressive projection networks for hashing.

3 PRELIMINARY

In this work, we consider cross-modal hash retrieval problems. The objective is to map samples from both modalities into the shared Hamming space, enabling efficient cross-modal retrieval. Specifically, let the visual vector space be $\mathcal{D}^v = \{\mathbf{f}_i^v\}_{i=1}^N$ and the text vector space be $\mathcal{D}^t = \{\mathbf{f}_i^t\}_{i=1}^N$, which is encoded by standard pre-trained visual and text encoders. We have N image-text pairs without label information. We employ neural networks $\phi^v(\cdot)$ and $\phi^t(\cdot)$ to map each visual and text feature vector \mathbf{f}_i^v and \mathbf{f}_i^t into the Hamming space as:

$$\mathbf{b}_i^v = \text{sign}(\phi^v(\mathbf{f}_i^v)), \quad \mathbf{b}_i^t = \text{sign}(\phi^t(\mathbf{f}_i^t)), \quad (1)$$

where \mathbf{b}_i^v and \mathbf{b}_i^t are L -length hash codes, *i.e.*, $\mathbf{b}_i^* \in \{-1, +1\}^L$, $* \in \{v, t\}$, and $\text{sign}(\cdot)$ is the sign function. The hash codes could be used for subsequent efficient retrieval. Therefore, we need to minimize the Hamming distances between semantically similar samples across modalities while maximizing the distances between dissimilar ones. The Hamming distance is calculated as $d(\mathbf{b}_i^*, \mathbf{b}_j^*) = \frac{1}{2}(L - \langle \mathbf{b}_i^*, \mathbf{b}_j^* \rangle)$, where L is the code length and $\langle \cdot, \cdot \rangle$ denotes the inner product.

4 METHODOLOGY

4.1 FRAMEWORK OVERVIEW

Sparse cross-modal connections in unsupervised scenarios pose challenges for modality alignment and cross-modal retrieval. The core idea of HINT is to establish a cross-modal encoding tree to recover the hierarchical structure across modalities and enhance the connections among local semantic clusters. Specifically, as illustrated in Figure 2, our HINT comprises three main components: ① *Hierarchical encoding tree construction*. Guided by the hierarchical structural entropy, we optimize the enhanced cross-modal relationship graph to obtain the encoding tree. ② *Cross-modal hash learning with modality mixup*. To minimize the heterogeneous gap between two modalities, HINT builds

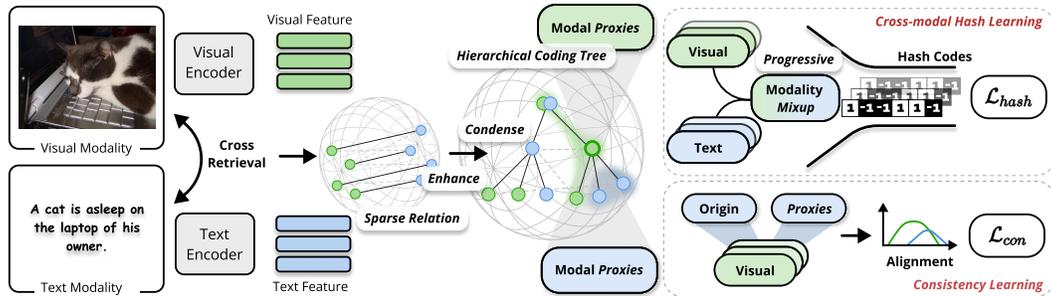


Figure 2: Overview of HINT, which constructs a hierarchical encoding tree from sparse cross-modal relationships. It then synthesizes modality proxies and performs progressive modality mixup and global-view consistency learning.

proxy samples for different modalities and progressively aligns them through a curriculum-based modality mixup mechanism (gradually transitioning from intra-modal to cross-modal alignment).
 ③ *Proxy-based consistency learning.* We optimize the distribution of cross-modal proxy samples, achieving a global-level alignment.

4.2 HIERARCHICAL CODING TREE CONSTRUCTION

To address the challenges posed by flat and sparse cross-modal connections, we construct a hierarchical encoding tree in an *enhance-and-condense* manner, as shown in Figure 3. First, we enhance the intra-modal connections within the relation graph. Then, guided by the structure entropy (Li et al., 2018; Zou et al., 2023), we condense the relation graph to obtain the hierarchical encoding tree. The encoding tree exhibits a hierarchical community structure, facilitating the hash learning.

Enhance. In unsupervised cross-modal retrieval, we primarily rely on cross-modal pairwise relations. We first construct an inter-modal relation graph $\mathcal{G}_{inter} = \{\mathcal{V}, \mathcal{E}_{inter}\}$, where $\mathcal{V} = \mathcal{D}^v \cup \mathcal{D}^t$ and $\mathcal{E}_{inter} = \{\mathbf{f}_i^v, \mathbf{f}_i^t\}_{i=1}^N$. Since cross-modal pairs only provide sparse supervision signals in unsupervised scenarios, we strategically employ KNN to enhance intra-modal relationships. This enables us to capture fine-grained local similarities and form cohesive bottom-level semantic communities, which serve as a robust foundation for subsequent hierarchical modeling. Since \mathcal{G}_{inter} is sparse and inadequate for cross-modal learning, we enrich the intra-modal relationships to construct tightly-knit low-level communities. Specifically, we turn to the cosine similarity within modality by $S_{(i,j)}^* = \cos(\mathbf{f}_i^*, \mathbf{f}_j^*)$, $* \in \{v, t\}$. We choose cosine similarity as it focuses on semantic directional alignment by normalizing vector magnitudes, which is crucial for cross-modal feature comparison. We then construct the intra-modal relationship graph $\mathcal{G}_{intra} = \{\mathcal{V}, \mathcal{E}_{intra}\}$ based on the similarity matrix by KNN manner:

$$\mathcal{E}_{intra} = \left\{ \left\{ \mathbf{f}_i^*, \mathbf{f}_j^* \right\} \mid j \in \text{topk}(\mathbf{f}_i^*, S, k) \right\}_{i=1}^N, \quad (2)$$

where k is set to 3 according to the hyperparameter study in Section 5.2. This choice balances information gain and noise robustness: larger k introduces noisy relationships, while smaller k fails to capture sufficient local semantics. Then, we merge the intra-modal relationships \mathcal{G}_{intra} and the inter-modal relationships \mathcal{G}_{inter} to obtain the cross-modal relationship graph $\mathcal{G}_{cross} = \{\mathcal{V}, \mathcal{E}_{intra} \cup \mathcal{E}_{inter}\}$.

Condense. To simplify the cross-modal semantics and construct the hierarchical relationship, we extract the encoding tree \mathcal{T} from \mathcal{G}_{cross} with the guidance of structural entropy, as illustrated in the center part of Figure 2. We first introduce the 1-D structural entropy as:

$$E^1(\mathcal{G}) = - \sum_{v \in \mathcal{V}} \frac{d_v}{\text{vol}(\mathcal{G})} \log \frac{d_v}{\text{vol}(\mathcal{G})}, \quad (3)$$

where d_v denotes the degree for node v and $\text{vol}(\cdot)$ denotes the sum of degree for nodes in \mathcal{G} . Then, we introduce the hierarchical structural entropy for the \mathcal{T} as:

$$E^{\mathcal{T}}(\mathcal{G}) = - \sum_{\alpha \in \mathcal{T}} \underbrace{\frac{g_\alpha}{\text{vol}(\mathcal{G})}}_{\text{information leakage}} \log \underbrace{\frac{\mathcal{V}_\alpha}{\mathcal{V}_{\alpha^-}}}_{\text{encoding efficiency}}. \quad (4)$$

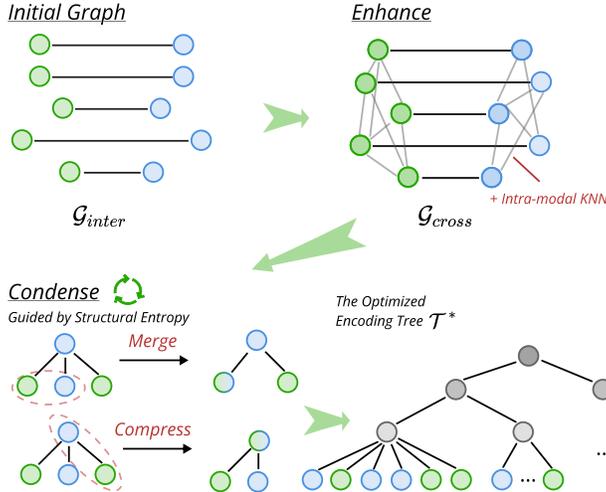


Figure 3: The construction pipeline of the encoding tree, involving graph enhancement and entropy-guided condensing operations.

Intuitively, the factor $\frac{g_\alpha}{\text{vol}(\mathcal{G})}$ captures the information leakage of community α , while the logarithmic term $\log \frac{\mathcal{V}_\alpha}{\mathcal{V}_{\alpha^-}}$ reflects the encoding efficiency of the hierarchy. In Eq. 4, \mathcal{G} is $\mathcal{G}_{\text{cross}}$, α is a node in \mathcal{T} , \mathcal{T}_α is the subtree rooted at α , \mathcal{T}_{α^-} is the subtree rooted at α 's parent node, g_α is the number of intra-modal relation links originating from the subtree \mathcal{T}_α , \mathcal{V}_α is the sum of degrees in subtree \mathcal{T}_α , and \mathcal{V}_{α^-} is the sum of degrees in subtree \mathcal{T}_{α^-} . Eq. 4 generalizes Eq. 3 to hierarchical communities and reduces to Eq. 3 when the hierarchy collapses to a flat partition.

Guided by structural entropy, we convert the cross-modal relation graph into an encoding tree through *Merge* and *Compress*, as shown in Figure 3. Firstly, we perform node *Merge* operation to generate a binary encoding tree. The node *Merge* operation merges nodes that belong to the same parent node. These nodes may be highly semantically similar, and merging them can reduce the overall structural entropy of the cross-modal encoding tree. The node *Merge* operation is defined as: $\mathcal{T}' = \text{Merge}_{\mathcal{T}}(\alpha, \beta)$, we check all nodes if it can introduce decrease in $E^{\mathcal{T}}(\mathcal{G})$.

Secondly, *Compress* operation is performed to optimize the encoding tree, mainly targeting adjacent nodes at different levels, and constructing local clusters. This is achieved by attempting to move the child encoding tree with α as the root to its parent node's parent node, thereby enabling compression of the cross-modal semantic graph. After encoding tree compression, if there are no child encoding trees connected to a parent node, this parent node can be contracted to its parent node. The *Compress* operation is defined as: $\mathcal{T}' = \text{Compress}_{\mathcal{T}}(\alpha, \beta)$. Similarly, we check the nodes and conduct the *Compress* operation if it can introduce entropy decrease. Overall, we optimize the cross-modal encoding tree following a greedy principle. Specifically, we traverse the tree nodes in a breadth-first search manner. We attempt the aforementioned operations, and if they can decrease the structural entropy, we execute the operation. While this greedy approach does not guarantee global optimality, it provides an efficient and practical solution that converges to stable structures with clear hierarchical communities, as validated by our experiments. The cross-modal encoding tree after optimization is defined as $\mathcal{T}^* = \arg \min (E^{\mathcal{T}}(\mathcal{G}))$, where \mathcal{T}^* is the optimized encoding tree. The resulting tree naturally captures semantic granularity transitions, with upper nodes representing broad categories, mid-level nodes capturing fine-grained concepts, and leaf nodes corresponding to specific instances. \mathcal{T}^* exhibits better cross-modal semantic properties. Specifically, they encompass more comprehensive local semantic motifs while mitigating the connections within high-density communities. These characteristics facilitate subsequent discriminative hash code learning and consistency learning.

4.3 STRUCTURE-GUIDED CROSS-MODAL HASH LEARNING

After obtaining the optimized cross-modal encoding tree \mathcal{T}^* , we use it for unsupervised hash learning. Compared to existing unsupervised cross-modal hashing works (Hu et al., 2022; Liu et al., 2017; Zhang et al., 2018; 2024b; Tu et al., 2023), our method can exploit local semantic communities, avoiding the bias caused by individual samples. Simultaneously, we jointly model both image and text modalities, mapping the vectors from distinct modalities into a unified binary space.

Proxy Construction. For each sample f_i^* , we sample its neighboring nodes $\mathcal{N}^+(f_i^*)$ with the same modality and $\mathcal{N}^-(f_i^*)$ with the opposite modality on the cross-modal encoding tree \mathcal{T}^* , and obtain the *proxy* samples via:

$$\begin{aligned} m_i^{same} &= \frac{1}{|\mathcal{N}^+(f_i^*)|} \sum_{j \in \mathcal{N}^+(f_i^*)} \phi^*(f_j^*), \\ m_i^{cross} &= \frac{1}{|\mathcal{N}^-(f_i^*)|} \sum_{j \in \mathcal{N}^-(f_i^*)} \phi^*(f_j^*), \end{aligned} \quad (5)$$

where f_i^* , $*$ $\in \{v, t\}$ is the vector, $\phi^*(\cdot)$ is the hash model we introduced in Eq. 1. The neighbor counts $|\mathcal{N}^+(f_i^*)|$ and $|\mathcal{N}^-(f_i^*)|$ are naturally determined by the tree structure \mathcal{T}^* and vary across nodes. Therefore, m_i^{same} for a text sample f_i^t (aggregating text neighbors) is distinct from m_i^{same} for a visual sample f_i^v (aggregating visual neighbors), as they are derived from different sets of modality-specific neighbors. A similar distinction applies to m_i^{cross} . The modal *proxy* samples can be viewed as a mediator between the two modalities, consisting of semantically similar nodes from the opposite modality, exhibiting better semantic robustness.

Modality Mixup. Leveraging *proxy* samples, we introduce a mixup mechanism (Oh et al., 2023; Verma et al., 2019; Huang et al., 2022) for progressively learning. Specifically, as Figure 4, m_i^{same} and m_i^{cross} are mixup to generate the hash codes:

$$\begin{aligned} b_i^{mix} &= \text{sign}\left(\frac{\lambda}{1+\lambda} m_i^{same} + \frac{1}{1+\lambda} m_i^{cross}\right), \\ \lambda &= \widehat{\text{MMD}}(\rho(m^{same}, \mathcal{B}), \rho(m^{cross}, \mathcal{B})), \end{aligned} \quad (6)$$

where λ is to measure the distribution difference of different modalities, and ρ is the cosine distance metrics, we sample in mini-batch \mathcal{B} and calculate $\widehat{\text{MMD}}$ using the Gaussian kernel (Long et al., 2015; Tolstikhin et al., 2016). Notably, λ is not a manually tuned hyperparameter but a data-driven metric computed via MMD, which dynamically quantifies the modality gap during training. This adaptive weighting mechanism implements a curriculum learning strategy (Hu et al., 2025). The theoretical foundation lies in manifold regularization (Verma et al., 2019): by interpolating in feature space between semantically purified proxy samples, we encourage smooth decision boundaries and force both modalities to find a shared semantic space where their interpolation remains meaningful for hash learning. As Figure 5, the evolution of λ empirically demonstrates this progressive alignment process.

Cross-modal Hash Learning. We employ the mixed hash code b_i^{mix} for hash learning. In the learning process, we sample in batches, and the other samples in the batch can serve as negative samples, enhancing the discriminative power of the hash codes. The objective function for cross-modal hash learning is:

$$\mathcal{L}_{hash} = - \sum_{i=1}^N \left(\log \frac{\exp(\langle f_i^*, b_i^{mix} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{B}|} \exp(\langle f_i^*, b_j^{mix} \rangle / \tau)} \right), \quad (7)$$

where $*$ $\in \{v, t\}$ and τ is the temperature parameter, which is set to 0.3 according to Section 5.2.

In HINT, we utilize the cross-modal encoding tree to guide hash learning. Since our cross-modal encoding tree has more comprehensive connections on local base groups, it can help align hash codes to more robust semantics. Meanwhile, the cross-modal encoding tree is sparser between communities, facilitating discriminability between groups and achieving discriminative hash codes.

Theoretical Discussion. This paragraph will discuss the limiting behavior of the cross-modal hashing loss \mathcal{L}_{hash} and demonstrate how it enables effective retrieval. In Eq. 7, the inner product $\langle f_i^*, b_i^{mix} \rangle$ serves as a similarity measure between the feature f_i^* and the hash code b_i^{mix} , and

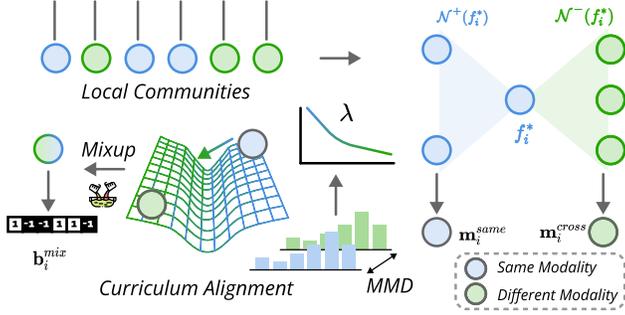


Figure 4: The Modality Mixup pipeline: generating modality-specific proxies, using their MMD to guide a curriculum-based mixup, and producing b_i^{mix} .

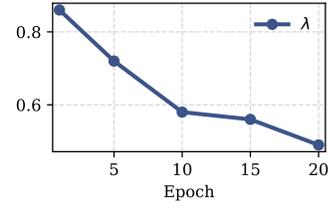


Figure 5: Evolution of λ during learning on MIRFlickr-25K.

the temperature τ controls the scale of the measured similarity. Since τ is relatively small in implementation, the following theorem follows Oh et al. (2023) and shows the convergence property of \mathcal{L}_{hash} .

Lemma 4.1. *For sufficiently large N and batch size $|\mathcal{B}|$, we have*

$$\lim_{\tau \rightarrow 0^+} \frac{1}{N} \mathcal{L}_{hash} \simeq \mathbb{E}[\|\mathbf{f}_i^v - \mathbf{b}_i^{mix}\|_2^2 + \|\mathbf{f}_i^t - \mathbf{b}_i^{mix}\|_2^2 - \min_{k \neq i} \|\mathbf{f}_i^v - \mathbf{b}_k^{mix}\|_2^2 - \min_{k \neq i} \|\mathbf{f}_i^t - \mathbf{b}_k^{mix}\|_2^2]. \quad (8)$$

The right-hand side of Eq. 8 can be used to measure the difference between the distances between positive pairs in comparison with the hardest negative pairs that are closest to the positive pair. Theorem 4.1 implies that minimizing \mathcal{L}_{hash} is equivalent to minimizing the triplet loss, and the smaller triplet loss implies more transferable representation and more effective retrieval. This theorem shows how \mathcal{L}_{hash} contributes to the cross-modal retrieval performance. The detailed proof is available in the Appendix B.

4.4 PROXY-BASED CONSISTENCY LEARNING

Due to the heterogeneous gap between modalities, it is necessary to introduce an additional global-view modal alignment mechanism to achieve better alignment and enhance the generalization ability.

Semantic Consistency Learning. We leverage the modal *proxy* samples for modality alignment learning, assuming that the original samples and their *proxy* samples should have similar semantic positions and similar distribution. Specifically, the modality semantics in a global view can be represented as the similarity distribution between samples and other samples within the same batch:

$$p(\mathbf{f}_i^*) = [\rho(\mathbf{f}_i^*, \mathbf{f}_j^*) \mid \mathbf{f}_j^* \in \mathcal{B}^-], \quad (9)$$

where \mathcal{B}^- includes the opposite modality instances within the same mini-batch, and $\rho(\cdot)$ is the cosine similarity function. Our objective is achieved by optimizing the KL divergence:

$$\mathcal{L}_{con} = \sum_{i=1}^{|\mathcal{B}|} (D_{KL}(p(\mathbf{f}_i^*) \parallel p(\mathbf{m}_i^{cross}))), \quad (10)$$

where $|\mathcal{B}|$ is the batch size, $p(\mathbf{f}_i^*)$ and $p(\mathbf{m}_i^{cross})$ are the semantic distributions of the i -th sample and its cross-modal *proxy* sample, respectively. By optimizing the consistency learning \mathcal{L}_{con} , we achieve modality alignment learning at a high level by leveraging the semantic-stable *proxy* samples.

Summary. Our method constructs a hierarchical encoding tree by unsupervised cross-modal mining and simultaneously leverages the encoding tree for cross-modal hash learning and semantic consistency learning. During the testing phase, the hierarchical encoding tree and proxy samples are not used. We directly employ the corresponding hash model to generate its hash code. This design ensures efficient retrieval with minimal computational overhead during inference time. Due to the non-differentiability of the $\text{sign}(\cdot)$ function, it is challenging to optimize the overall objective. Therefore, we adopt $\tanh(\cdot)$ as a surrogate during the training process. The whole algorithm is summarized in Algorithm 1 and Appendix A. The computational complexity and time efficiency are discussed in the Appendix C.

5 EXPERIMENT

5.1 EXPERIMENTAL SETTINGS

🔹 **Datasets.** In order to comprehensively evaluate the performance of HINT, we conduct experiments on three popular public datasets: MIRFlickr-25K (Huiskes & Lew, 2008), NUS-WIDE (Rasiwasia et al., 2010), and MS-COCO (Lin et al., 2014). The detailed information is available in Appendix E.2. 🔹 **Baselines.** We compare our method HINT with 15 baselines from related fields, including the latest state-of-the-art works. Details in Appendix E.1. 🔹 **Implementation Details.** To ensure a fair comparison, we implement our method based on the latest SOTA works (Zhang et al., 2024b; Hu et al., 2022). Details in Appendix E.3.

5.2 RESULTS

Hamming Ranking. The experiments on cross-modal retrieval benchmarks show that the proposed HINT consistently outperforms baseline approaches across different code lengths (16-128 bits). Key

Table 1: Comparison of MAP performance (%) across various cross-modal hashing methods.

Methods	MIRFlickr-25K				NUS-WIDE				MS-COCO			
	16	32	64	128	16	32	64	128	16	32	64	128
<i>Image → Text</i>												
CVH	62.0	60.8	59.4	58.3	48.7	49.5	45.6	41.9	50.3	50.4	47.1	42.5
LSSH	59.7	60.9	60.6	60.5	44.2	45.7	45.0	45.1	48.4	52.5	54.2	55.1
CMFH	55.7	55.7	55.6	55.7	33.9	33.8	34.3	33.9	36.6	36.9	37.0	36.5
FSH	58.1	61.2	63.5	66.2	55.7	56.5	59.8	63.5	53.9	54.9	57.6	58.7
MTFH	50.7	51.2	55.8	55.4	29.7	29.7	27.2	32.8	39.9	29.3	29.5	39.5
FOMH	57.5	64.0	69.1	65.9	30.5	30.5	30.6	31.4	37.8	51.4	57.1	60.1
DCH	59.6	60.2	62.6	63.6	39.2	42.2	43.0	43.6	42.2	42.0	44.6	46.8
DGCPN	65.1	68.3	71.8	72.4	60.1	61.8	63.1	64.0	55.6	56.9	57.8	58.0
UCHSTM	70.1	71.5	72.4	72.3	62.5	63.5	64.6	64.4	55.8	57.2	57.6	57.3
UCCH	71.6	72.6	72.8	73.2	62.1	62.3	64.0	64.5	56.0	56.2	56.6	57.4
UDDH	71.4	72.9	74.0	74.6	63.7	64.2	65.1	65.9	56.8	57.8	59.0	59.9
HuggingHash+	71.6	73.2	74.3	74.5	63.9	64.8	65.6	66.4	57.1	58.3	59.4	60.5
DEMO	71.8	73.3	73.4	74.3	64.6	64.8	66.2	66.4	57.5	57.8	58.6	60.5
GCRH	71.0	72.2	72.7	73.3	63.9	64.0	65.3	65.5	56.8	57.0	57.8	59.6
VTM-UCH	71.9	73.6	73.9	74.5	64.6	65.1	66.0	66.6	57.6	58.2	58.8	60.3
HINT	72.9	74.4	75.1	75.5	65.1	65.5	66.5	67.3	58.5	59.5	60.4	61.1
<i>Text → Image</i>												
CVH	62.9	61.5	59.9	58.7	47.0	47.5	44.4	41.2	50.6	50.8	48.6	42.9
LSSH	60.2	59.8	59.8	59.7	47.3	48.2	47.1	45.7	49.0	52.2	54.7	56.0
CMFH	55.3	55.3	55.3	55.3	30.6	30.6	30.6	30.6	34.6	34.6	34.6	34.6
FSH	57.6	60.7	63.5	66.0	56.9	60.4	65.1	66.6	53.7	52.4	56.4	57.3
MTFH	51.4	52.4	51.8	58.1	35.3	31.4	39.9	41.0	33.5	37.4	30.0	33.4
FOMH	58.5	64.8	71.9	68.8	30.2	30.4	30.0	30.6	36.8	48.4	55.9	59.5
DCH	61.2	62.3	65.3	66.5	37.9	43.2	44.4	45.9	42.1	42.8	45.4	47.1
DGCPN	65.3	68.2	71.2	71.5	60.5	62.6	63.7	64.4	55.0	56.6	57.8	57.7
UCHSTM	69.5	71.1	71.3	72.3	63.2	64.3	65.1	65.2	55.5	56.7	57.8	57.3
UCCH	70.3	71.2	72.0	72.1	62.5	63.7	65.0	65.2	56.4	57.3	57.2	58.1
UDDH	70.5	71.6	72.8	73.5	64.5	65.2	66.0	66.6	56.6	57.5	58.5	59.4
HuggingHash+	70.7	72.0	73.2	73.8	64.8	65.7	66.5	67.0	56.9	57.9	59.0	60.1
DEMO	70.8	71.9	72.2	72.8	65.4	65.5	66.9	67.1	57.2	57.9	58.3	59.7
GCRH	70.1	71.0	71.3	71.8	64.5	64.7	66.0	66.2	56.5	57.3	57.4	58.8
VTM-UCH	71.4	72.1	72.4	73.0	65.6	65.7	66.9	67.3	57.4	58.0	58.3	59.6
HINT	72.0	73.1	74.0	74.6	66.0	66.6	67.3	67.8	58.2	59.0	59.8	60.8

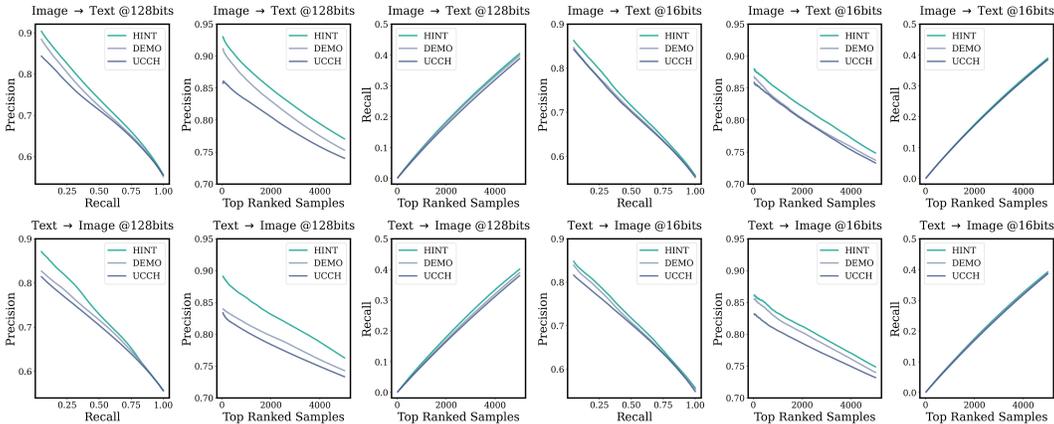


Figure 6: Hash lookup performance with 128-bit and 16-bit codes on the MIRFlickr-25K dataset. The Precision-Recall curves, Precision-N curves, and Recall-N curves are shown from left to right.

findings show that: ❶ Deep unsupervised hashing approaches typically outperform traditional ones, ❷ Supervised methods struggle when labeled data is limited, ❸ The method shows improved performance on challenging sub-tasks like Text→Image retrieval, and ❹ The hierarchical modeling approach proves more effective than other deep cross-modal methods for achieving progressive cross-modal alignment. Additionally, HINT demonstrates strong robustness against noisy data, maintain-

Table 2: Ablation studies. The component columns ‘‘KNN, Tree, Curr, Con’’ respectively denote intra-modal KNN, hierarchical encoding tree, curriculum-based mixup, and proxy-based learning.

Methods	Components				MIRF-25K		NUS-WIDE		MS-COCO	
	KNN	Tree	Curr	Con	I→T	T→I	I→T	T→I	I→T	T→I
HINT V1					73.2	72.0	64.0	65.1	57.9	58.5
HINT V2	✓				73.8	72.8	65.2	65.7	59.1	58.9
HINT V3	✓	✓			74.2	73.6	65.9	66.4	60.0	59.5
HINT V4	✓	✓	✓		75.1	74.1	67.0	67.3	60.7	60.2
HINT	✓	✓	✓	✓	75.5	74.6	67.3	67.8	61.1	60.8

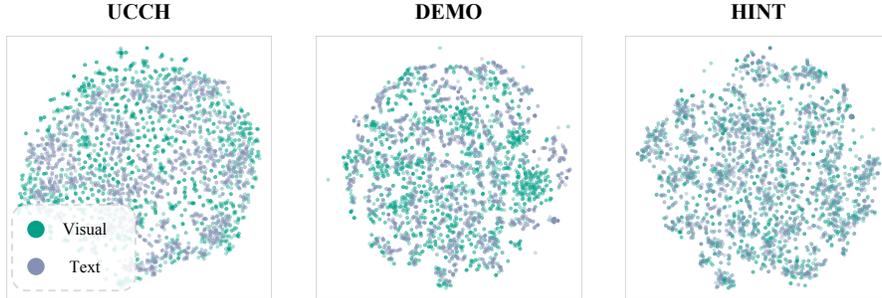


Figure 7: The t-SNE projection of hash codes from different modalities. Among competing methods, HINT shows the best ability of modal alignment.

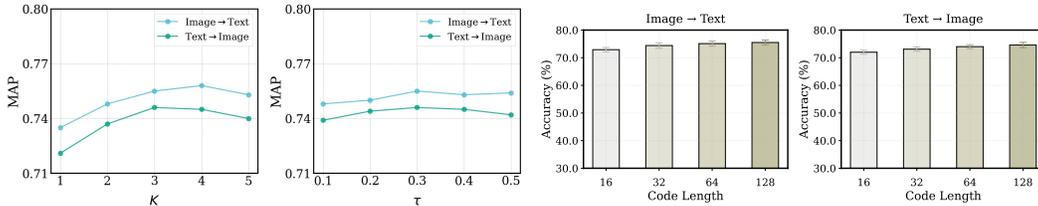
ing superior performance even with 10% corrupted pairs. Detailed noise robustness analysis and results are provided in Appendix C.6. As shown in Table 1, performance generally improves with increasing bit length as longer hash codes provide larger Hamming space for encoding more information, though with diminishing returns at higher lengths. The more pronounced improvements on Text→Image tasks can be attributed to text feature sparsity and lower initial quality compared to visual features. HINT’s hierarchical structure and proxy-based neighborhood aggregation provide more substantial benefits for the inherently more challenging tasks.

Hash Lookup. To comprehensively analyze HINT’s performance, we evaluate Precision-Recall, Precision-N, and Recall-N curves with 128-bit and 64-bit codes on MirFlickr-25K. From the results in Figure 6, HINT performs better in comparison to these baselines across all metrics, aligning with the MAP scores from Hamming ranking. In summary, HINT exhibited optimal performance in cross-modal hash retrieval. Experiments for other code lengths are in Appendix C.1.

Visualization. We provide a t-SNE visualization analysis of HINT’s performance using 128-bit hash codes on the MirFlickr-25K dataset, comparing with DEMO and UCCH, distinguishing different modalities with distinct colors. As shown in Figure 7, HINT demonstrates a superior ability to map representations from different modalities into a unified hash space, exhibiting higher alignment between text and visual modalities. The visualization suggests that HINT effectively aligns modalities and learns hash codes with generalization capabilities.

Ablation Study. We compare the following variants of HINT: V1, which only uses text-image pairs without consistency learning (\mathcal{L}_{con}), where \mathbf{b}_{mix}^* is obtained solely from the opposite modality; V2, which uses both text-image pairs and intra-modal KNN without \mathcal{L}_{con} , where \mathbf{b}_{mix}^* is obtained by averaging opposite modality and KNN samples; V3, which constructs the hierarchical encoding tree with sample selection but without curriculum-based progressive alignment and without consistency loss \mathcal{L}_{con} ; V4, which uses \mathcal{L}_{hash} consistent with the full model but excludes \mathcal{L}_{con} . As shown in Table 2, the full model achieves optimal performance, with hierarchical encoding tree and progressive alignment (V3 and V4) yielding the most improvements. We also explored iterative tree updates but found static construction provides better efficiency-performance trade-off, with detailed analysis in Appendix C.8. Additional experiments comparing different similarity metrics for tree construction demonstrate cosine similarity’s superiority over L1/L2 distances, with detailed analysis in Appendix C.7. Additional ablation studies are available in Appendix C.2.

Sensitivity Analysis. We analyze the hyperparameters K and τ . As shown in Figure 8a, increasing K from 1 to 3 improves performance on both retrieval tasks, validating the benefits of enhanced



(a) Sensitive analysis of K and τ . HINT demonstrates robustness to hyperparameters. (b) Stability analysis of HINT. Error bars indicate standard deviation.

Figure 8: Sensitivity and stability analyses on the MIRFlickr-25K dataset.

cross-modal relationships. However, further increasing K to 5 introduces noisy relationships and decreases performance. Our experiments demonstrate HINT’s remarkable stability, with performance fluctuation remaining within a narrow 2% margin when varying τ from 0.1 to 0.5. The model consistently outperforms baselines across most parameter settings. Based on these findings, we set $K = 3$ and $\tau = 0.3$ as the default values in our experiments.

Stability Analysis. We conducted 5 independent runs with different seeds. As shown in Figure 8b, the results show that HINT exhibits remarkable stability, with performance variations consistently remaining below 1% standard deviation across different code lengths. Details in Appendix C.9.

Time Efficiency. We provide a comprehensive efficiency analysis covering model complexity, training cost, and retrieval speed. As shown in Table 5, HINT achieves superior MAP performance with training efficiency comparable to DEMO and UDDH, demonstrating an excellent efficiency-performance trade-off. For retrieval efficiency, HINT and all compared hashing methods use identical Hamming distance computation on binary codes, thus sharing the same retrieval speed. Table 4 demonstrates that hash-based retrieval achieves $26\text{-}31\times$ speedup over dense vector approaches. Details in Appendix C.5 and C.4.

6 CONCLUSION

This paper investigates the problem of efficient cross-modal retrieval through cross-modal hashing. We propose HINT that leverages hierarchical structural entropy to guide the construction of a cross-modal encoding tree, which has tightly connected local clusters. By incorporating progressive mixup for proxy-based alignment and consistency learning from a global perspective, we enhance the generalization capability of the produced hash codes. In the end, we demonstrate the effectiveness of our HINT via comprehensive experiments on benchmark datasets.

ETHICS STATEMENT

Our research adheres to the ICLR Code of Ethics. All datasets used in this study are publicly available. The code and related materials will be appropriately released to ensure transparency and reproducibility of our work.

ACKNOWLEDGMENT

Ming Zhang and Junyu Luo are supported by grants from the National Natural Science Foundation of China (NSFC Grant Number 62276002).

REFERENCES

- Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Qiang Yang. Transitive hashing network for heterogeneous multimedia retrieval. In *Proc. of Association for the Advancement of Artificial Intelligence*, volume 31, 2017.
- Dapeng Chen, Min Wang, Haobin Chen, Lin Wu, Jing Qin, and Wei Peng. Cross-modal retrieval with heterogeneous graph embedding. In *Proc. of ACM International Conference on Multimedia*, pp. 3291–3300, 2022.

- Haonan Chen, Zhicheng Dou, Kelong Mao, Jiongnan Liu, and Ziliang Zhao. Generalizing conversational dense retrieval via LLM-cognition data augmentation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024.
- Kan Chen, Trung Bui, Chen Fang, Zhaowen Wang, and Ram Nevatia. Amc: Attention guided multi-modal correlation learning for image search. In *Proc. of Computer Vision and Pattern Recognition*, 2017.
- Yu Chen, Lingfei Wu, and Mohammed Zaki. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. *Advances in neural information processing systems*, 33:19314–19326, 2020.
- Wanqing Cui, Keping Bi, Jiafeng Guo, and Xueqi Cheng. MORE: Multi-mOdal REtrieval augmented generative commonsense reasoning. In *Findings of the Association for Computational Linguistics: ACL 2024*, 2024.
- Guiguang Ding, Yuchen Guo, Jile Zhou, and Yue Gao. Large-scale cross-modality search via collective matrix factorization hashing. *IEEE Transactions on Image Processing*, 2016a.
- Kun Ding, Bin Fan, Chunlei Huo, Shiming Xiang, and Chunhong Pan. Cross-modal hashing via rank-order preserving. *IEEE Transactions on Multimedia*, 2016b.
- Haozhi Fan and Yuan Cao. Vision-guided text mining for unsupervised cross-modal hashing with community similarity quantization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 2843–2851, 2025.
- Zijun Gao, Jun Wang, Guoxian Yu, Zhongmin Yan, Carlotta Domeniconi, and Jinglin Zhang. Long-tail cross modal hashing. In *Proc. of Association for the Advancement of Artificial Intelligence*, 2023.
- Xuri Ge, Fuhai Chen, Joemon M Jose, Zhilong Ji, Zhongqin Wu, and Xiao Liu. Structured multi-modal feature embedding and alignment for image-sentence retrieval. In *Proc. of ACM International Conference on Multimedia*, pp. 5185–5193, 2021.
- Xuri Ge, Songpei Xu, Fuhai Chen, Jie Wang, Guoxin Wang, Shan An, and Joemon M Jose. 3shnet: Boosting image-sentence retrieval via visual semantic-spatial self-highlighting. *Information Processing & Management*, 61(4):103716, 2024.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- Peng Hu, Hongyuan Zhu, Jie Lin, Dezhong Peng, Yin-Ping Zhao, and Xi Peng. Unsupervised contrastive cross-modal hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- Xiaoxing Hu, Kaicheng Yang, Ziyong Feng, Qi Ming, Zonghao Guo, Xiang An, Junchi Yan, and Xue Yang. Proclip: Progressive vision-language alignment via llm-based embedder. *arXiv preprint arXiv:2510.18795*, 2025.
- Hailang Huang, Zhijie Nie, Ziqiao Wang, and Ziyu Shang. Cross-modal and uni-modal soft-label alignment for image-text retrieval. In *Proc. of Association for the Advancement of Artificial Intelligence*, volume 38, pp. 18298–18306, 2024.
- Zhenyu Huang, Guocheng Niu, Xiao Liu, Wenbiao Ding, Xinyan Xiao, Hua Wu, and Xi Peng. Learning with noisy correspondence for cross-modal matching. *Proc. of Advances in Neural Information Processing Systems*, 2021.
- Zhipeng Huang, Jiawei Liu, Liang Li, Kecheng Zheng, and Zheng-Jun Zha. Modality-adaptive mixup and invariant decomposition for rgb-infrared person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pp. 1034–1042, 2022.
- Mark J Huiskes and Michael S Lew. The mir flickr retrieval evaluation. In *Proc. of ACM international conference on Multimedia information retrieval*, 2008.

- Qing-Yuan Jiang and Wu-Jun Li. Discrete latent factor model for cross-modal hashing. *IEEE Transactions on Image Processing*, 2019.
- Weike Jin, Zhou Zhao, Pengcheng Zhang, Jieming Zhu, Xiuqiang He, and Yueting Zhuang. Hierarchical cross-modal graph consistency learning for video-text retrieval. In *Proceedings of the 44th International ACM SIGIR Conference on research and development in information retrieval*, pp. 1114–1124, 2021.
- Benno Krojer, Vaibhav Adlakha, Vibhav Vineet, Yash Goyal, Edoardo Ponti, and Siva Reddy. Image retrieval from contextual descriptions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- Shaishav Kumar and Raghavendra Udupa. Learning hash functions for cross-view similarity search. In *Proc. of International Joint Conference on Artificial Intelligence*, 2011.
- Saehyung Lee, Sangwon Yu, Junsung Park, Jihun Yi, and Sungroh Yoon. Interactive text-to-image retrieval with large language models: A plug-and-play approach. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024.
- Angsheng Li and Yicheng Pan. Structural information and dynamical complexity of networks. *IEEE Transactions on Information Theory*, 2016.
- Angsheng Li, Xianchen Yin, Bingxiang Xu, Danyang Wang, Jimin Han, Yi Wei, Yun Deng, Ying Xiong, and Zhihua Zhang. Decoding topologically associating domains with ultra-low resolution hi-c data by graph structural entropy. *Nature communications*, 2018.
- Chao Li, Cheng Deng, Lei Wang, De Xie, and Xianglong Liu. Coupled cyclegan: Unsupervised hashing network for cross-modal retrieval. In *Proc. of Association for the Advancement of Artificial Intelligence*, 2019.
- Fengling Li, Bowen Wang, Lei Zhu, Jingjing Li, Zheng Zhang, and Xiaojun Chang. Cross-domain transfer hashing for efficient cross-modal retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024a.
- Wenyan Li, Jiaang Li, Rita Ramos, Raphael Tang, and Desmond Elliott. Understanding retrieval robustness for retrieval-augmented image captioning. *arXiv preprint arXiv:2406.02265*, 2024b.
- Yunxin Li, Baotian Hu, Yuxin Ding, Lin Ma, and Min Zhang. A neural divide-and-conquer reasoning framework for image retrieval from linguistically complex text. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023.
- Zhi Li, Yifan Liu, and Yin Zhang. Back-modality: Leveraging modal transformation for data augmentation. *Proc. of Advances in Neural Information Processing Systems*, 2024c.
- Meiyu Liang, Junping Du, Zhengyang Liang, Yongwang Xing, Wei Huang, and Zhe Xue. Self-supervised multi-modal knowledge graph contrastive hashing for cross-modal search. In *Proc. of Association for the Advancement of Artificial Intelligence*, volume 38, pp. 13744–13753, 2024.
- Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Proc. of Advances in Neural Information Processing Systems*, 2022.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proc. of European Conference on Computer Vision*, 2014.
- Hong Liu, Rongrong Ji, Yongjian Wu, Feiyue Huang, and Baochang Zhang. Cross-modality binary code learning via fusion similarity hashing. In *Proc. of Computer Vision and Pattern Recognition*, 2017.
- Wei Liu, Jun Wang, Rongrong Ji, Yu-Gang Jiang, and Shih-Fu Chang. Supervised hashing with kernels. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2074–2081. IEEE, 2012.

- Xin Liu, Zhikai Hu, Haibin Ling, and Yiu-ming Cheung. Mtfh: A matrix tri-factorization hashing framework for efficient cross-modal retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019a.
- Xuanwu Liu, Guoxian Yu, Carlotta Domeniconi, Jun Wang, Yazhou Ren, and Maozu Guo. Ranking-based deep cross-modal hashing. In *Proc. of Association for the Advancement of Artificial Intelligence*, 2019b.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proc. of International Conference on Machine Learning*. PMLR, 2015.
- Xu Lu, Lei Zhu, Zhiyong Cheng, Jingjing Li, Xiushan Nie, and Huaxiang Zhang. Flexible online multi-modal hashing for large-scale multimedia retrieval. In *Proc. of ACM International Conference on Multimedia*, 2019.
- Junyu Luo, Yusheng Zhao, Xiao Luo, Zhiping Xiao, Wei Ju, Li Shen, Dacheng Tao, and Ming Zhang. Cross-domain diffusion with progressive alignment for efficient adaptive retrieval. *IEEE Transactions on Image Processing*, 2025.
- Xiao Luo, Haixin Wang, Daqing Wu, Chong Chen, Minghua Deng, Jianqiang Huang, and Xian-Sheng Hua. A survey on deep hashing methods. *ACM Transactions on Knowledge Discovery from Data*, 2023.
- Zeyu Ma, Siwei Wang, Xiao Luo, Zhonghui Gu, Chong Chen, Jinxing Li, Xian-Sheng Hua, and Guangming Lu. Harr: Learning discriminative and high-quality hash codes for image retrieval. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(5):1–23, 2024.
- Georgii Mikriukov, Mahdyar Ravanbakhsh, and Begüm Demir. Unsupervised contrastive hashing for cross-modal retrieval in remote sensing. In *Proc. of International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 2022.
- Changdae Oh, Junhyuk So, Hoyoon Byun, YongTaek Lim, Minchul Shin, Jong-June Jeon, and Kyungwoo Song. Geodesic multi-modal mixup for robust fine-tuning. In *Proc. of Advances in Neural Information Processing Systems*, 2023.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. of International Conference on Machine Learning*, 2021.
- Nikhil Rasiwasia, Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Gert RG Lanckriet, Roger Levy, and Nuno Vasconcelos. A new approach to cross-modal multimedia retrieval. In *Proc. of ACM International Conference on Multimedia*, 2010.
- Fumin Shen, Chunhua Shen, Wei Liu, and Heng Tao Shen. Supervised discrete hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 37–45, 2015.
- Xiaobo Shen, Yinfan Chen, Weiwei Liu, Yuhui Zheng, Quan-Sen Sun, and Shirui Pan. Graph convolutional multi-label hashing for cross-modal retrieval. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- Mingyang Song, Liping Jing, and Yi Feng. Match more, extract better! hybrid matching model for open domain web keyphrase extraction. In *Findings of the Association for Computational Linguistics ACL 2024*, pp. 17–27, 2024.
- Yuan Sun, Zhenwen Ren, Peng Hu, Dezhong Peng, and Xu Wang. Hierarchical consensus hashing for cross-modal retrieval. *IEEE Transactions on Multimedia*, 26:824–836, 2023.
- Yuan Sun, Jian Dai, Zhenwen Ren, Yingke Chen, Dezhong Peng, and Peng Hu. Dual self-paced cross-modal hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 15184–15192, 2024.

- Ilya O Tolstikhin, Bharath K Sriperumbudur, and Bernhard Schölkopf. Minimax estimation of maximum mean discrepancy with radial kernels. *Proc. of Advances in Neural Information Processing Systems*, 2016.
- Rong-Cheng Tu, Xian-Ling Mao, Qinghong Lin, Wenjin Ji, Weize Qin, Wei Wei, and Heyan Huang. Unsupervised cross-modal hashing via semantic text mining. *IEEE Transactions on Multimedia*, 2023.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *International conference on machine learning*, pp. 6438–6447. PMLR, 2019.
- Dongsheng Wang, Miaoge Li, Xinyang Liu, MingSheng Xu, Bo Chen, and Hanwang Zhang. Tuning multi-mode token-level prompt alignment across modalities. *Proc. of Advances in Neural Information Processing Systems*, 2024a.
- Haixin Wang, Hao Wu, Jinan Sun, Shikun Zhang, Chong Chen, Xian-Sheng Hua, and Xiao Luo. Idea: An invariant perspective for efficient domain adaptive image retrieval. *Advances in Neural Information Processing Systems*, 36:57256–57275, 2023.
- Jinpeng Wang, Ziyun Zeng, Bin Chen, Yuting Wang, Dongliang Liao, Gongfu Li, Yiru Wang, and Shu-Tao Xia. Hugs bring double benefits: Unsupervised cross-modal hashing with multi-granularity aligned transformers. *International Journal of Computer Vision*, pp. 1–33, 2024b.
- Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. Semi-supervised hashing for scalable image retrieval. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 3424–3431. IEEE, 2010.
- Rukai Wei, Yu Liu, Heng Cui, Yanzhao Xie, and Ke Zhou. Graph contrastive-and-reconstructive hashing for unsupervised cross-modal retrieval. *Data Science and Engineering*, pp. 1–17, 2025.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proc. of Computer Vision and Pattern Recognition*, 2018.
- Xing Xu, Fumin Shen, Yang Yang, Heng Tao Shen, and Xuelong Li. Learning discriminative binary codes for large-scale cross-modal retrieval. *IEEE Transactions on Image Processing*, 2017.
- Chenggang Yan, Biao Gong, Yuxuan Wei, and Yue Gao. Deep multi-view enhancement hashing for image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Jun Yu, Hao Zhou, Yibing Zhan, and Dacheng Tao. Deep graph-neighbor coherence preserving network for unsupervised cross-modal hashing. In *Proc. of Association for the Advancement of Artificial Intelligence*, 2021.
- Bin Zhang, Yue Zhang, Junyu Li, Jiazhou Chen, Tatsuya Akutsu, Yiu-Ming Cheung, and Hongmin Cai. Unsupervised dual deep hashing with semantic-index and content-code for cross-modal retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024a.
- Donglin Zhang, Zhiwen Wang, Zhikai Hu, and Xiao-Jun Wu. Htvr: Hierarchical text-to-video retrieval based on relative similarity. *Pattern Recognition*, pp. 112145, 2025.
- Fan Zhang, Xian-Sheng Hua, Chong Chen, and Xiao Luo. DEMO: A Statistical Perspective for Efficient Image-Text Matching. *arXiv preprint arXiv:2405.11496*, 2024b.
- Jian Zhang, Yuxin Peng, and Mingkuan Yuan. Unsupervised generative adversarial cross-modal hashing. In *Proc. of Association for the Advancement of Artificial Intelligence*, 2018.
- Liangli Zhen, Peng Hu, Xu Wang, and Dezhong Peng. Deep supervised cross-modal retrieval. In *Proc. of Computer Vision and Pattern Recognition*, 2019.
- Jile Zhou, Guiguang Ding, and Yuchen Guo. Latent semantic sparse hashing for cross-modal similarity search. In *Proc. of ACM SIGIR conference on Research & development in information retrieval*, 2014.

Lei Zhu, Chaoqun Zheng, Weili Guan, Jingjing Li, Yang Yang, and Heng Tao Shen. Multi-modal hashing for efficient multimedia retrieval: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 2023.

Dongcheng Zou, Hao Peng, Xiang Huang, Renyu Yang, Jianxin Li, Jia Wu, Chunyang Liu, and Philip S Yu. Se-gsl: A general and effective graph structure learning framework through structural entropy optimization. In *Proc. of ACM Web Conference*, 2023.

A ALGORITHM

We present the optimization algorithm of our method in Algorithm 1. The algorithm first constructs a cross-modal relationship graph and optimizes the hierarchical encoding tree. Then during training, it performs modality mixup and consistency learning to generate effective hash codes through back-propagation.

Algorithm 1 Optimization Algorithm of HINT

Require: Visual modality \mathcal{D}^v ; Text modality \mathcal{D}^t ;

Ensure: Hashing model $\phi^v(\cdot)$ and $\phi^t(\cdot)$;

- 1: Construct the cross-modal relationship graph \mathcal{G}_{cross} ;
 - 2: Optimize the hierarchical encoding tree \mathcal{T}^* ;
 - 3: **for** each epoch **do**
 - 4: **for** each batch **do**
 - 5: Sample $\mathcal{B}^v, \mathcal{B}^t$ from $\mathcal{D}^v, \mathcal{D}^t$;
 - 6: Construct proxy samples $\mathbf{m}^{same}, \mathbf{m}^{cross}$ using Eq. 5;
 - 7: Perform modality mixup and generate hash code \mathbf{b}^{mix} with \mathcal{T}^* by Eq. 6;
 - 8: Calculate the \mathcal{L}_{hash} and \mathcal{L}_{con} ;
 - 9: Update parameters through back-propagation;
 - 10: **end for**
 - 11: **end for**
-

B PROOF OF LEMMA 4.1

Proof of Lemma 4.1. Denote \mathcal{N}_v and \mathcal{N}_t denote the index sets of visual vector and text vector, correspondingly. We assume the batch number $|\mathcal{B}|$ is equal to N to reduce the notation burden and by definition,

$$\mathcal{L}_{hash} = - \sum_{i \in \mathcal{N}_v} \left(\log \frac{\exp(\langle \mathbf{f}_i^*, \mathbf{b}_i^{mix} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{N}|} \exp(\langle \mathbf{f}_i^*, \mathbf{b}_j^{mix} \rangle / \tau)} \right) - \sum_{i \in \mathcal{N}_t} \left(\log \frac{\exp(\langle \mathbf{f}_i^*, \mathbf{b}_i^{mix} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{N}|} \exp(\langle \mathbf{f}_i^*, \mathbf{b}_j^{mix} \rangle / \tau)} \right). \quad (11)$$

We focus on the first term of the right-hand side of the above equation. By the proof of Proposition D.2. in Oh et al. (2023),

$$\begin{aligned} & \lim_{\tau \rightarrow 0^+} -\frac{1}{N} \sum_{i \in \mathcal{N}_v} \left(\log \frac{\exp(\langle \mathbf{f}_i^*, \mathbf{b}_i^{mix} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{N}|} \exp(\langle \mathbf{f}_i^*, \mathbf{b}_j^{mix} \rangle / \tau)} \right) \\ &= \lim_{\tau \rightarrow 0^+} \frac{1}{N} \sum_{i \in \mathcal{N}_v} \frac{1}{\tau} \max\{\max_j \{\langle \mathbf{f}_i^*, \mathbf{b}_j^{mix} - \mathbf{b}_i^{mix} \rangle\}, 0\}. \end{aligned} \quad (12)$$

Since \mathbf{m}_i^{same} and \mathbf{m}_i^{cross} are continuous random variables, then

$$\mathbb{P}(\|\mathbf{b}_i^{mix}\|^2 = L) = \mathbb{P}\left(\left\| \text{sign}\left(\frac{\lambda}{1+\lambda} \mathbf{m}_i^{same} + \frac{1}{1+\lambda} \mathbf{m}_i^{cross}\right)\right\|^2 = L\right) = 1. \quad (13)$$

and

$$\begin{aligned} & \|\mathbf{f}_i^v - \mathbf{b}_i^{mix}\|_2^2 - \min_{j \neq i} \|\mathbf{f}_i^v - \mathbf{b}_j^{mix}\|_2^2 \\ &= \max_{j \neq i} \|\mathbf{f}_i^v\|_2^2 + \|\mathbf{b}_i^{mix}\|_2^2 - 2\langle \mathbf{f}_i^v, \mathbf{b}_i^{mix} \rangle - \|\mathbf{f}_i^v\|_2^2 - \|\mathbf{b}_j^{mix}\|_2^2 + 2\langle \mathbf{f}_i^v, \mathbf{b}_j^{mix} \rangle \\ &= 2 \max_{j \neq i} \langle \mathbf{f}_i^v, \mathbf{b}_j^{mix} - \mathbf{b}_i^{mix} \rangle, \end{aligned} \quad (14)$$

which implies that

$$\lim_{\tau \rightarrow 0^+} -\frac{1}{N} \sum_{i \in \mathcal{N}_v} \left(\log \frac{\exp(\langle \mathbf{f}_i^*, \mathbf{b}_i^{mix} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{N}|} \exp(\langle \mathbf{f}_i^*, \mathbf{b}_j^{mix} \rangle / \tau)} \right) \simeq \|\mathbf{f}_i^v - \mathbf{b}_i^{mix}\|_2^2 - \min_{j \neq i} \|\mathbf{f}_i^v - \mathbf{b}_j^{mix}\|_2^2. \quad (15)$$

□

Table 3: Additional ablation studies with different code lengths.

Methods	MIRFlickr-25K							
	Image→Text				Text→Image			
	16bit	32bit	64bit	128bit	16bit	32bit	64bit	128bit
HINT V1	71.5	72.4	72.6	73.2	70.5	71.1	71.7	72.0
HINT V2	71.8	72.8	73.0	73.7	70.8	71.4	72.2	72.8
HINT V3	72.0	73.4	73.8	74.2	71.5	72.4	73.2	73.6
HINT V4	72.6	73.9	74.8	75.2	71.5	72.8	73.9	74.3
Full Model	72.9	74.4	75.1	75.5	72.0	73.1	74.0	74.9

Table 4: Retrieval time cost (ms) varies with code length.

	16 Bit	32 Bit	48 Bit	64 Bit	96 Bit	128 Bit
Hash Code	16.7	18.0	19.4	19.9	21.8	22.2
Dense Vector	441.4	491.0	543.0	602.3	657.7	696.6
Speed Up	26.5×	27.2×	28.0×	30.2×	30.1×	31.4×

C ADDITIONAL EXPERIMENTAL RESULTS

C.1 ADDITIONAL EXPERIMENTS FOR HASH LOOKUP

To comprehensively evaluate the performance of the proposed HINT, we present precision-recall curves, precision-N curves, and recall-N curves on the MirFlickr-25K dataset with code lengths of 32 and 64 bits. As shown in Figure 10, 9, our method consistently outperforms other approaches, which is consistent with the corresponding MAP score based on Hamming ranking. Additionally, we compute the precision and recall rates of the top-N retrieved results, demonstrating HINT’s persistent advantage. Our method achieves superior performance in cross-modal hash retrieval.

C.2 ADDITIONAL ABLATION EXPERIMENTS

To evaluate the effectiveness of these components, we introduce four variants:

- HINT V1, which only employs image-text pairs for cross-modal hash learning without the global-view consistency learning module;
- HINT V2, which combines image-text pairs with intra-modal KNN for cross-modal hash learning, also without the consistency learning module;
- HINT V3, which uses the average of two modalities as the learning target and builds a hierarchical encoding tree with sample selection, but without curriculum-based progressive modal alignment (Equation 6), also without the global-view consistency learning module;
- HINT V4, which adopts the full model’s Lhash but excludes consistency learning.

As shown in Table 3, our results demonstrate that our complete HINT achieves optimal performance, confirming the importance of each component. Furthermore, our ablation study reveals that hierarchical encoding trees and progressive alignment yield significant improvements (V2 and V3), validating our motivation.

C.3 COMPUTATIONAL COMPLEXITY

We mainly discuss the computational complexity of the additionally introduced encoding tree construction. Assuming the number of data points is N (including samples from different modalities), the time cost for calculating the similarity matrix is $O(N^2)$, the cost for constructing the KNN graph is $O(N)$, and the cost for optimizing the cross-modal encoding tree is $O(N \log^2 N)$ (Li & Pan, 2016). It is worth noting that the most time-consuming similarity matrix calculation can be

Table 5: Model complexity and training efficiency comparison on MIRFlickr-25K with 128-bit.

Method	Params (M)	Total Time (min)	MAP (%)
UCHSTM	52.5	182	72.3
UCCH	88.0	58	73.2
UDDH	28.9	45	74.6
HuggingHash+	40.1	211	74.5
DEMO	86.2	59	74.3
HINT	84.8	61	75.5

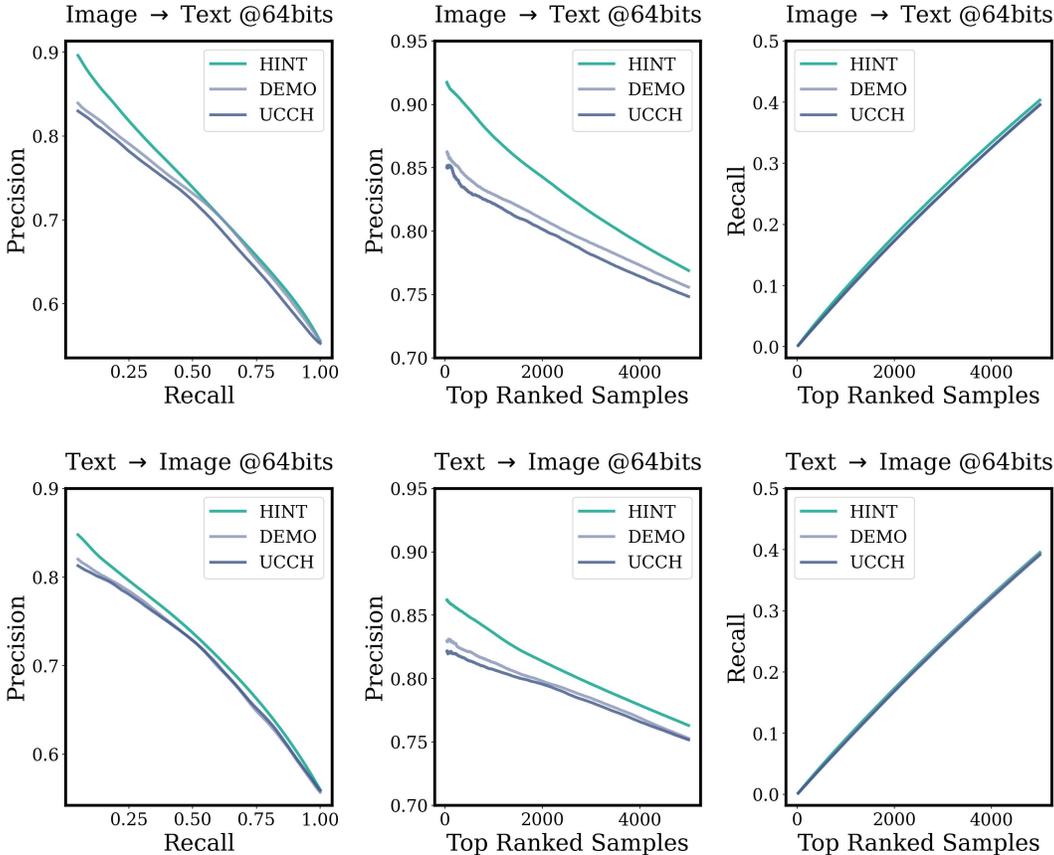


Figure 9: Hash lookup performance with 64 bits codes on the MIRFlickr-25K dataset. The precision-recall curves, precision-N curves, and recall-N curves are shown from left to right.

accelerated by parallel computing. Moreover, the encoding tree construction is only performed once at the beginning of the training, so the additional computational complexity and time consumption brought to the overall training process are negligible.

C.4 TIME EFFICIENCY

We conducted a speed test between HINT and dense vector retrieval on an Intel Xeon CPU E5-2697 v4 (2.30GHz), as illustrated in Table 4. The speed test is conducted on the MIRFlickr-25K dataset with a retrieval database of 10^5 items. We perform 10^3 runs and report the average retrieval speed cost (ms). Hash methods excel in enabling efficient and scalable image retrieval, especially for large-scale datasets, due to fast Hamming distance computation. In contrast, existing pre-trained models only output dense vectors, resulting in slower computation. Table 4 compares the efficiency of hash codes and dense codes generated by our model and a pre-trained model at various bit lengths. The

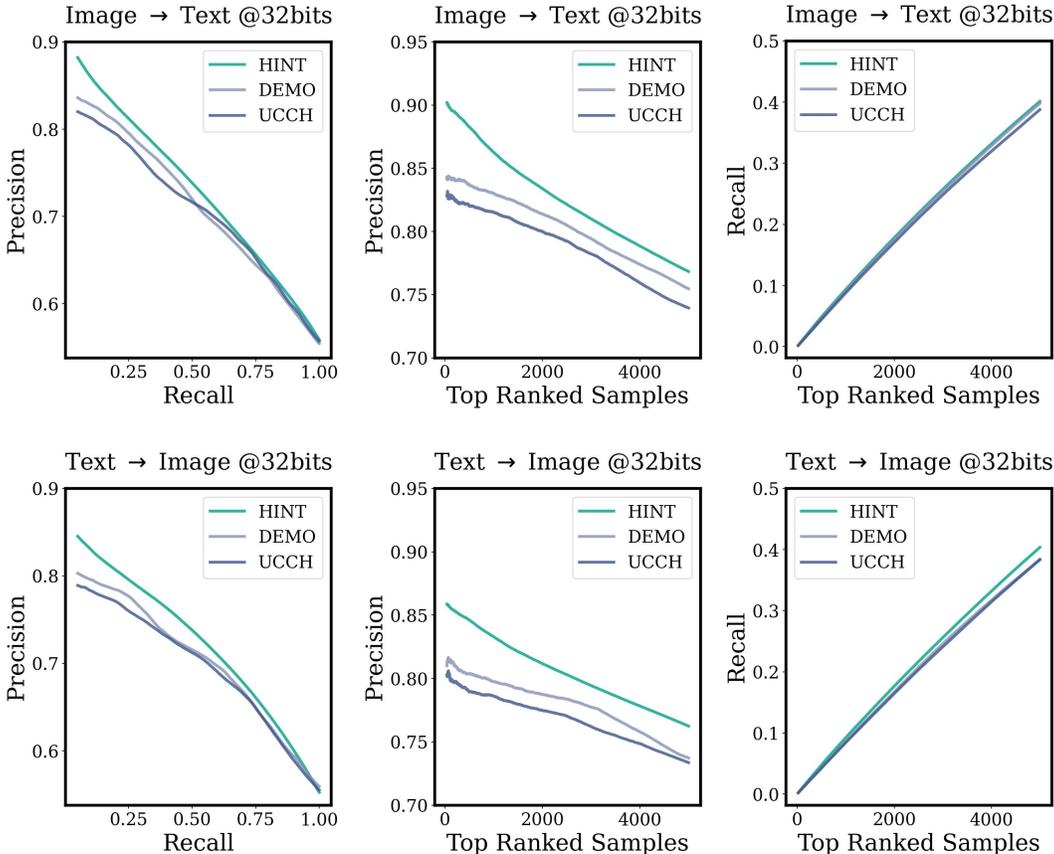


Figure 10: Hash lookup performance with 32 bits codes on the MIRFlickr-25K dataset. The precision-recall curves, precision-N curves, and recall-N curves are shown from left to right.

results clearly demonstrate that hash codes achieve significantly faster retrieval speeds than deep feature codes, confirming their superiority, particularly in large-scale image retrieval scenarios.

C.5 COMPUTATIONAL EFFICIENCY ANALYSIS

While HINT introduces additional computational steps through the hierarchical encoding tree, we have implemented several optimizations to ensure practical efficiency:

- Controlled Time Complexity:** The hierarchical encoding tree construction has a complexity of $O(N \log^2 N)$ and is executed only once during the initial training phase. Our experiments on MIRFlickr-25K show that tree construction takes <3 minutes (single GPU), representing <5% of total training time. Compared to existing methods like DEMO and UCCH, HINT does not significantly increase the overall training duration.
- Memory-Efficient Design:** The encoding tree is stored using compressed relation triplets (parent-child-edge weight) instead of maintaining complete similarity matrices. Furthermore, the hash code generation phase is completely decoupled from the encoding tree, eliminating the need to load tree structures during inference and conserving deployment resources.
- Practical Scalability:** The encoding tree’s one-time construction and reusability make it particularly suitable for large-scale applications. This design choice significantly amortizes the initial computational investment across multiple training sessions.

These results demonstrate that HINT achieves superior performance while maintaining competitive training efficiency through its optimized design. The comprehensive model complexity and training efficiency comparison with existing methods is shown in Table 5.

Table 6: Noise robustness comparison on MIRFlickr-25K (128 bit) with 10% corrupted pairs.

Method	I→T	T→I	I→T(10% <i>n</i>)	T→I(10% <i>n</i>)
UCCH	73.2	73.2	70.7	72.6
DEMO	74.3	74.3	71.4	73.3
HINT	75.5	73.8	74.4	72.9

Table 7: Performance comparison of different similarity metrics on MIRFlickr-25K.

Metric	I→T	T→I
Cosine	75.5	74.6
L2	74.8	74.2
L1	73.5	72.8

C.6 NOISE ROBUSTNESS ANALYSIS

To evaluate HINT’s robustness against noisy data, we conducted experiments by randomly corrupting 10% of text-image pairs in the MIRFlickr-25K dataset. The results demonstrate HINT’s superior noise resilience through both architectural design and experimental validation:

Architectural Robustness: HINT’s hierarchical encoding tree provides two-level noise adaptation:

- The tree construction process inherently suppresses individual outliers by aggregating local semantic communities. Proxy samples, generated through neighbor feature averaging, effectively smooth out the impact of noisy samples within local communities.
- The cross-modal consistency learning module (L_{con}) constrains the influence of outliers on hash space mapping by enforcing semantic distribution alignment between proxy and original samples from a global perspective.

As shown in Table 6, HINT consistently maintains higher performance under noisy conditions, with minimal degradation compared to baseline methods. This demonstrates that the encoding tree’s hierarchical structure effectively identifies and mitigates the interference of mismatched pairs in cross-modal alignment.

C.7 ADDITIONAL ANALYSIS OF SIMILARITY METRICS

We conducted experiments comparing different similarity metrics for tree construction on MIRFlickr-25K. The results demonstrate that cosine similarity achieves optimal performance due to three key advantages:

- **Directional Consistency:** Cosine similarity focuses on semantic directional alignment by normalizing vector magnitudes
- **Loss Function Alignment:** The training objective relies on inner product similarity, which aligns with cosine similarity computation
- **Feature Space Compatibility:** Hamming distance is not suitable since features are not yet binarized during tree construction

The results show cosine similarity’s 1.5-2.0% performance advantage over alternative metrics, confirming that feature directional alignment is more crucial than absolute distance for tree construction.

C.8 ITERATIVE TREE CONSTRUCTION ANALYSIS

Iterative structural optimization is a promising direction (Chen et al., 2020). Our experiments reveal that a static tree construction strategy achieves better balance between efficiency and performance:

- **Empirical Results:** We compared HINT with HINT-ITER (dynamic tree updates every 5 epochs) across multiple datasets:

Table 8: Performance comparison between static and iterative tree construction.

Method	MIRFlickr-25K		MS-COCO		Training Time (h)
	I→T	T→I	I→T	T→I	
HINT	75.5	74.6	61.1	60.8	2.1
HINT-ITER	75.0	74.7	60.5	60.4	2.9

Table 9: Stability analysis of HINT across different code lengths on MIRFlickr-25K dataset. Results show mean MAP scores \pm standard deviation over five runs.

Task	32 bits	64 bits	96 bits	128 bits
I→T	72.9 \pm 0.81	74.4 \pm 0.99	75.1 \pm 0.92	75.5 \pm 0.88
T→I	72.0 \pm 0.75	73.1 \pm 0.75	74.0 \pm 0.72	74.6 \pm 0.96

The results show that HINT-ITER achieves comparable but slightly lower performance while requiring more training time. Two fundamental reasons explain this phenomenon:

- **Robust Initial Structure:** Our one-time tree construction leverages hierarchical structural entropy to recover semantically coherent communities, providing a stable foundation for proxy sample generation. Iterative refinement struggles to further improve this already optimized structure.
- **Stability-Aware Alignment:** The curriculum-based mixup mechanism and consistency learning rely on stable neighborhood relationships to progressively align modalities. Frequent tree updates disrupt this process, similar to how unstable negative samples degrade contrastive learning (He et al., 2020).

While our current approach suits existing tasks, we acknowledge potential benefits of dynamic structures for specific scenarios (e.g., evolving data streams), which we leave for future work.

C.9 STABILITY ANALYSIS

To rigorously assess the stability of HINT, we conducted extensive experiments with five independent runs using different random seeds. Table 9 presents the mean performance and standard deviations across different code lengths for both Image-to-Text (I→T) and Text-to-Image (T→I) retrieval tasks on the dataset.

The results demonstrate that HINT maintains consistent performance with remarkably low variance across different code lengths. The standard deviations consistently remain below 1% for both retrieval directions, indicating strong robustness to random initialization. This stability can be attributed to our hierarchical encoding tree structure and curriculum-based progressive alignment strategy, which provide reliable guidance for hash code learning regardless of initialization conditions.

C.10 TAIL CLASS PERFORMANCE ANALYSIS

To evaluate HINT’s robustness under class imbalance, we conducted tail class analysis on MIRFlickr-25K dataset. We categorized the 24 classes into three groups based on sample counts: Head (8 classes), Medium (10 classes), and Tail (6 classes). As shown in Figure 11, HINT demonstrates consistent advantages across all class groups, with particularly significant improvements on Tail classes. This superior performance can be attributed to HINT’s hierarchical structure: the encoding tree leverages cross-modal connections to augment sparse intra-modal neighborhoods for tail classes, while the hierarchical community structure ensures these connections remain semantically coherent rather than introducing noise from head classes.

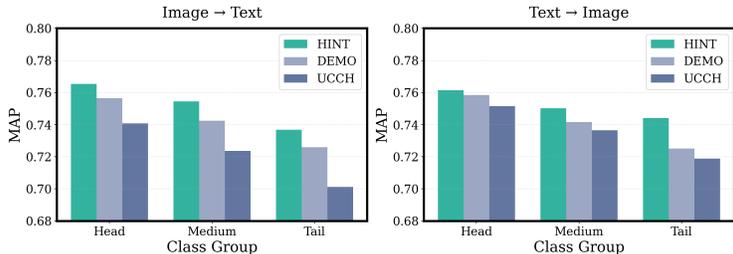


Figure 11: Performance comparison on Head, Medium, and Tail class groups on MIRFlickr-25K (128-bit). HINT demonstrates consistent advantages across all groups, with more significant improvements on Tail classes compared to baselines.

D ADDITIONAL DISCUSSION

D.1 RATIONALE FOR HIERARCHICAL MINING OF RELATIONSHIPS

Hierarchical semantic structures are inherent in real-world data. Visual and textual content naturally form multi-level conceptual taxonomies. For instance, a general category like “Objects” can be decomposed into “Animals” and “Vehicles”. “Animals” can be further subdivided into “Domestic” and “Wild”, with “Domestic” containing specific instances like “Cats” and “Dogs” (e.g., “Maine Coon”, “Siamese”). Similarly, “Vehicles” might branch into “Ground” and “Air” transport, with “Ground” including “Cars” and “Trucks”.

Prevailing unsupervised cross-modal hashing methods often rely on flat representations of image-text pair relationships. Such flat structures exhibit limitations in capturing these intrinsic hierarchical dependencies. Specifically, they may:

- Treat the semantic dissimilarity between disparate pairs (e.g., “Cat-Dog” vs. “Cat-Car”) undifferentiatedly, failing to recognize varying degrees of relatedness based on hierarchical proximity.
- Implicitly assume transitive relationships, an assumption that does not consistently hold for complex semantic relationships across different levels of abstraction.
- Struggle to ensure that instances within a sub-category (e.g., “Maine Coon” and “Siamese” under “Cats”) are represented as being semantically closer to each other than to instances from distant categories (e.g., “Cars”).

Our proposed HINT addresses these limitations through the Hierarchical Encoding Tree, which explicitly discovers and models inherent semantic hierarchies by optimizing structural entropy. This hierarchical approach offers several advantages:

- It facilitates a deeper exploration of semantic relationships that extend beyond direct, observed pairings, uncovering latent community structures.
- It enables the generation of more generalizable hash codes that are grounded in these discovered semantic communities, rather than isolated instances.
- It promotes smoother and more effective cross-modal alignment through the use of proxy samples derived from semantic neighborhoods and a curriculum learning strategy, thereby more effectively bridging the semantic gap between modalities.

By capturing these multi-level containment relationships, HINT can learn hash codes that better reflect real-world semantic structures, leading to improved retrieval performance. Current flat modeling approaches, by contrast, which treat “Cat-Dog” and “Cat-Car” similarity differences with equal weight, miss this crucial hierarchical semantic information, significantly impeding their ability to learn generalizable hash codes that align with complex real-world semantic organizations.

E ADDITIONAL IMPLEMENTATION DETAILS

E.1 BASELINE DETAILS

The introduction of the baseline methods is as follows:

- *CVH* (Kumar & Udupa, 2011) transforms the learning process into a manageable feature-based hashing problem.
- *LSSH* (Zhou et al., 2014) is an effective iterative strategy, which explores the correlations between multi-modal representations and bridges the semantic gaps in the latent semantic space.
- *CMFH* (Ding et al., 2016a) exploits cross-modal decomposition to establish strong connections.
- *FSH* (Liu et al., 2017) alternates optimization to learn consistent binary codes and hash functions.
- *MTFH* (Liu et al., 2019a) produces modality-specific binary codes with varying lengths while ensuring the comparability of heterogeneous data.
- *FOMH* (Lu et al., 2019) integrates a self-weighted and flexible multi-modal fusion strategy, enabling robust fusion even when missing modalities.
- *DCH* (Xu et al., 2017) utilizes an efficient optimization algorithm to produce the optimized bit recursively.
- *DGCPN* (Yu et al., 2021) is a graph neighborhood approach, which explore the relationships between data points and their neighbors through a graph neighborhood approach, thereby enhancing the accuracy of data similarity measurement.
- *UCHSTM* (Tu et al., 2023) employs a custom-designed similarity loss to rectify any ill-defined similarities in the instance similarity matrix.
- *UCCH* (Hu et al., 2022) uses contrastive learning to enforce alignment between different modalities and unified binary representations, focusing on leveraging discriminative information from all pairs.
- *UDDH* (Zhang et al., 2024a) is a dual deep hashing architecture that combines semantic indexing with content codes for cross-modal retrieval.
- *HuggingHash+* (Wang et al., 2024b) is a transformer-based multi-granularity learning framework for unsupervised cross-modal hashing.
- *DEMO* (Zhang et al., 2024b) utilizes multi-view augmentation to represent each image, followed by parameterized distribution divergence to ensure robust similarity structures.
- *GCRH* (Wei et al., 2025) is a contrastive-and-reconstructive framework combining global graph contrastive learning with local graph reconstruction to reduce the modality gap.
- *VTM-UCH* (Fan & Cao, 2025) utilizes vision-guided text mining by combining vision-language models with community detection for refining hash code distribution.

E.2 DATASET DETAILS

We conduct experiments on three widely used public datasets:

- *MIRFlickr-25K* (Huiskes & Lew, 2008) contains 25,000 text-image pairs. Each text sample is embedded using the Bag-of-Words (BoW) strategy.
- *NUS-WIDE* (Rasiwasia et al., 2010) comprises 269,498 text-image pairs with multiple labels from 81 categories, where each text is embedded into a 1000-dimensional vector via BoW.
- *MS-COCO* (Lin et al., 2014) includes 123,287 text-image pairs with multiple labels from 80 categories. Each text sample is embedded into a 2026-dimensional BoW vector.

Following the problem settings of the latest baseline (Zhang et al., 2024b), each dataset is divided into a query set and a retrieval set. During the training process, only text-image pair information is accessible without label information.

E.3 IMPLEMENTATION DETAILS

Our method is implemented based on the latest baseline (Zhang et al., 2024b). For training, we maintain consistency with baseline approaches by selecting 10,000 samples as our training set. We utilize pre-extracted visual and text feature vectors, which are mapped to the Hamming space through two-layer MLPs with a dimension of 512. The implementation is done using PyTorch framework, with all experiments conducted on a single NVIDIA A40 GPU.

E.4 EVALUATION METRICS

Mean Average Precision (MAP) is a comprehensive metric widely used to evaluate retrieval performance in cross-modal hashing research (Wang et al., 2010; Liu et al., 2012; Shen et al., 2015). The MAP score has a range of 0 to 1, where higher values indicate better retrieval performance. It works by calculating the average precision for each query, followed by taking the mean across the queries in the test set. This provides a single-figure measure that reflects system performance across all relevant documents and all recall levels. MAP considers both precision and recall aspects of the retrieval system, making it particularly suitable for evaluating hashing-based retrieval systems where we are concerned with the overall ranking quality of results.

F THE USE OF LARGE LANGUAGE MODELS

In this article, we use LLM for language polishing and retrieval of the latest research works. We confirm that we take full responsibility for the contents written in this paper.