# IBEC-MVC: INFORMATION BOTTLENECK ENHANCED CONTRASTIVE MULTI-VIEW CLUSTERING VIA PROGRESSIVE MATRIX FACTORIZATION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Multi-view clustering has emerged as a fundamental problem in machine learning, aiming to discover hidden patterns by leveraging complementary information across multiple data representations. However, existing methods suffer from information loss during dimensionality reduction and fail to effectively balance view-specific characteristics with cross-view consistency. In this paper, we propose IBEC-MVC (Information Bottleneck Enhanced Contrastive Multi-View Clustering), a novel framework that integrates three key innovations: (1) Information Bottleneck-guided Progressive Matrix Factorization (IB-PMF) that systematically compresses representations while preserving clustering-relevant information through theoretically grounded constraints, (2) Entropy-Regularized Self-Supervised Contrastive Learning (ER-SCL) that enhances discriminative structure in latent representations while preventing cluster collapse, and (3) Mutual Information Maximization for View Consistency (MIM-VC) that optimally fuses multi-view information through principled information-theoretic objectives. Our theoretical analysis provides information preservation guarantees for the progressive factorization process, while extensive experiments on benchmark datasets demonstrate significant improvements over state-of-the-art methods.

## 1 INTRODUCTION

Multi-view clustering discovers hidden patterns by leveraging complementary information across diverse data representations, with broad applications spanning social media analysis, biomedical research, computer vision, and natural language processing Fang et al. (2023); Chowdhury et al. (2025). Different views capture distinct yet complementary aspects of the underlying data structure, and their principled integration yields more robust clustering than single-view approaches. In social networks, user behaviors manifest through textual content, social connections, and temporal patterns; in biomedicine, patient data encompasses genomic sequences, clinical measurements, and imaging modalities—where multi-view clustering has demonstrated remarkable effectiveness in disease subtyping Liu et al. (2024b); Zhao et al. (2024).

Traditional approaches relied on matrix factorization and kernel methods to find low-dimensional representations preserving cross-view clustering structure Kumar et al. (2023), but struggled with scalability and heterogeneous data integration. Deep learning introduced architectures that learn complex nonlinear mappings and hierarchical representations Yang et al. (2024); Sun et al. (2024), and contrastive paradigms maximize agreement between positive pairs while minimizing negative-pair similarity across views Li et al. (2024); Huang et al. (2024). Current methods nevertheless face fundamental limitations in information preservation, discriminative learning, and optimal view fusion.

Despite these advances, three fundamental challenges remain: ***Challenge C1: Information Loss During Dimensionality Reduction.*** High-dimensional multi-view feature spaces pose computational and statistical challenges Zhou et al. (2024). Traditional dimensionality reduction discards clustering-relevant information, and existing matrix factorization approaches lack principled mechanisms for choosing what to preserve, resulting in arbitrary loss of critical discriminative features Brown et al. (2024); Zhang et al. (2023). ***Challenge C2: View Inconsistency and Discrimi-***



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

***native Structure Learning.*** Multi-view data contains both consistent cross-view and view-specific discriminative information Diallo et al. (2023). Current contrastive approaches suffer from cluster collapse, where samples map to similar representations and discriminative structure is lost Xiao et al. (2024), and struggle to balance view-specific characteristics against cross-view consistency, yielding over-smoothed or inconsistent results Yuan et al. (2024). ***Challenge C3: Lack of Principled View Fusion Mechanisms.*** Heterogeneous views with varying dimensionalities, statistics, and noise characteristics complicate unified integration Chen et al. (2024b). Existing fusion approaches treat all views equally or rely on heuristic weighting without theoretical foundations, under-using complementary information Liu et al. (2023).

To address these challenges, we propose **IBEC-MVC** (Information Bottleneck Enhanced Contrastive Multi-View Clustering), a theoretically grounded framework providing: (1) **Information Bottleneck-guided Progressive Matrix Factorization (IB-PMF)** that employs information-theoretic principles to compress high-dimensional representations while guaranteeing clustering-relevant information preservation (*addressing C1*); (2) **Entropy-Regularized Self-Supervised Contrastive Learning (ER-SCL)** that enhances discriminative structure while preventing cluster collapse through entropy constraints balancing view-specific information with cross-view consistency (*addressing C2*); (3) **Mutual Information Maximization for View Consistency (MIM-VC)** that optimally fuses multi-view information by maximizing mutual information between view-specific and fused representations (*addressing C3*). Extensive experiments demonstrate state-of-the-art performance with significant improvements across benchmark datasets.

The main contributions are:

- We propose **IB-PMF**, a theoretically grounded dimensionality reduction approach preserving clustering-relevant information through information-theoretic constraints, providing compression guarantees.
- We develop **ER-SCL**, a contrastive learning framework maintaining discriminative structure while ensuring cross-view consistency through entropy regularization, preventing cluster collapse.
- We introduce **MIM-VC**, a principled view fusion mechanism adaptively combining multi-view information by maximizing mutual information, ensuring proportional view contributions.
- We conduct extensive experiments on five benchmark datasets, achieving average improvements over state-of-the-art methods, with comprehensive ablation studies validating each component.

## 2 Related Work

**Multi-View Clustering**. Multi-view clustering leverages complementary information across data representations to achieve superior clustering performance. Recent advances focus on deep learning frameworks that integrate heterogeneous views while preserving discriminative structures. Li et al. categorize deep multi-view clustering into autoencoder-based, generative, and contrastive frameworks Li et al. (2025b), and the shift from shallow methods to deep architectures has yielded substantial gains on complex multi-view data Moujahid et al. (2025); Zhang et al. (2024a).

**Non-Negative Matrix Factorization**. Non-negative matrix factorization (NMF) is a fundamental technique for clustering and dimensionality reduction, valued for its interpretability and parts-based representations. Recent advances focus on robust variants that tolerate noisy and corrupted data Zhang et al. (2024b); Wang et al. (2024a), including robust self-supervised symmetric NMF that improves clustering stability under challenging data conditions Liu et al. (2024a); Chen et al. (2024a).

## 3 Methodology

As illustrated in Figure 1, our framework integrates three innovations: Information Bottleneck-guided Progressive Matrix Factorization (IB-PMF), Entropy-Regularized Self-Supervised Contrastive Learning (ER-SCL), and Mutual Information Maximization for View Consistency (MIM-VC). IB-PMF applies three-layer progressive factorization with information bottleneck control to compress features while preserving clustering information; ER-SCL enhances discriminability via

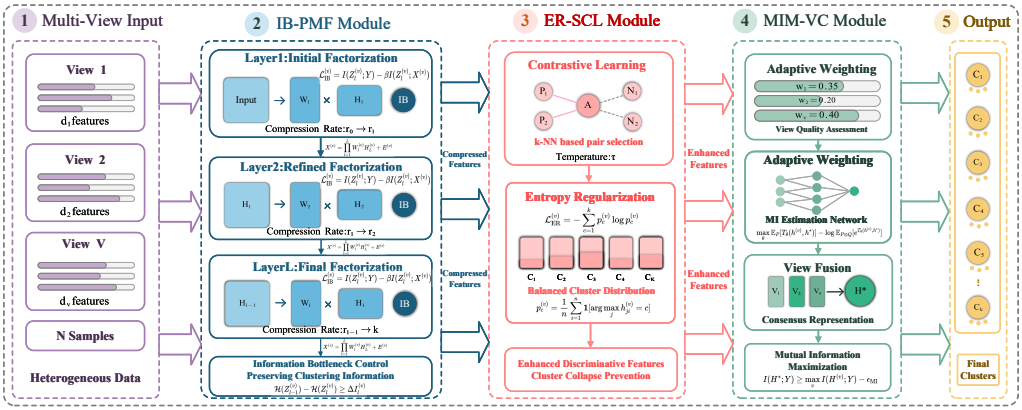


Figure 1: Overview of our proposed multi-view clustering framework.

contrastive learning with $k$-NN pair selection and prevents cluster collapse through entropy regularization; and MIM-VC fuses views via adaptive weighting and quality assessment to produce a consensus representation with balanced cluster distributions.

### 3.1 Problem Formulation

Given $V$ views $\{\mathbf{X}^{(v)}\}_{v=1}^{V}$ with $\mathbf{X}^{(v)} \in \mathbb{R}^{d_v \times n}$, our goal is to learn a unified clustering assignment matrix $\mathbf{H} \in \mathbb{R}^{k \times n}$ capturing the intrinsic cluster structure across all views, optimizing information flow while preserving view-specific characteristics and ensuring cross-view consistency.

### 3.2 Information Bottleneck-guided Progressive Matrix Factorization (IB-PMF)

Traditional matrix factorization methods for multi-view clustering lose information during dimensionality reduction, often discarding clustering-relevant signals along with noise. We therefore introduce the Information Bottleneck (IB) principle into progressive matrix factorization, providing a principled way to compress representations while preserving task-relevant information. The core insight is that effective clustering representations should maximize mutual information with cluster labels $\mathbf{Y}$ while minimizing it with irrelevant variations in the original data. For each view $v$ and layer $l$, we formulate the IB objective as:

$$\mathcal{L}_{IB}^{(v)} = I(\mathbf{Z}_l^{(v)}; \mathbf{Y}) - \beta I(\mathbf{Z}_l^{(v)}; \mathbf{X}^{(v)}) \tag{1}$$

Where $\mathbf{Z}_l^{(v)}$ denotes the $l$-th layer representation for view $v$, and $\beta$ controls the trade-off between information compression and preservation.

We design a progressive factorization scheme Where information is gradually compressed through multiple layers while maintaining clustering-relevant features. The factorization at layer $l$ follows:

$$\mathbf{X}^{(v)} = \mathbf{W}_1^{(v)}\mathbf{W}_2^{(v)}\cdots\mathbf{W}_L^{(v)}\mathbf{H}_L^{(v)} + \mathbf{E}^{(v)} = \prod_{l=1}^{L}\mathbf{W}_l^{(v)}\mathbf{H}_L^{(v)} + \mathbf{E}^{(v)} \tag{2}$$

Where $\mathbf{W}_l^{(v)} \in \mathbb{R}^{r_{l-1}\times r_l}$ are the progressive factor matrices with decreasing ranks $r_0 > r_1 > \cdots > r_L$, and $\mathbf{E}^{(v)}$ represents the reconstruction error.

Each layer is constrained by an information-theoretic requirement that ensures controlled information compression. Building upon the IB objective in Eq. equation 1, we derive the layer-wise constraint:

$$H(\mathbf{Z}_{l-1}^{(v)}) - H(\mathbf{Z}_l^{(v)}) \geq \Delta I_l^{(v)} \tag{3}$$

Where $H(\cdot)$ denotes entropy and $\Delta I_l^{(v)}$ is the minimum information compression at layer $l$. This constraint guarantees that each layer removes at least $\Delta I_l^{(v)}$ bits of information, enabling systematic noise removal while preserving clustering structure.

From the progressive factorization in Eq. equation 2 and the entropy constraint in Eq. equation 3, we can establish the following information preservation guarantee:

**Lemma 3.1** (Information Preservation in IB-PMF)**.** *For each layer $l$ in the IB-PMF framework, the clustering-relevant information is preserved with the following guarantee:*

$$I(\mathbf{Z}_l^{(v)}; \mathbf{Y}) \geq (1 - \beta^l) I(\mathbf{X}^{(v)}; \mathbf{Y}) \tag{4}$$

Proof can see Appendix A.1. To balance information preservation and compression efficiency, we employ a rate-distortion framework. Building upon the progressive factorization in Eq. equation 2 and incorporating the information preservation guarantee from Eq. equation 4, the optimization objective for each layer incorporates both reconstruction accuracy and information rate:

$$\min_{\mathbf{W}_l^{(v)}, \mathbf{H}_l^{(v)}} \|\mathbf{Z}_{l-1}^{(v)} - \mathbf{W}_l^{(v)} \mathbf{H}_l^{(v)}\|_F^2 + \lambda_l R(\mathbf{H}_l^{(v)}) \tag{5}$$

Where $R(\mathbf{H}_l^{(v)}) = -\sum_{i,j} h_{l,ij}^{(v)} \log h_{l,ij}^{(v)}$ is the rate term encouraging sparsity and interpretable representations.

### 3.3 Entropy-Regularized Self-Supervised Contrastive Learning (ER-SCL)

While IB-PMF provides information-theoretically sound representations, it may not fully capture the discriminative structure needed for clustering. We therefore introduce a self-supervised contrastive learning mechanism operating in each view's latent space. The key insight is that samples in the same cluster should have similar cross-view representations, while samples from different clusters should be distinguishable; unlike supervised contrastive learning, positive and negative pairs must be constructed without ground-truth labels.

**Contrastive Learning in Latent Space**. For each sample $i$ in view $v$, we construct positive pairs $\mathcal{P}_i$ based on $k$-nearest neighbors in the original feature space, under the assumption that nearby samples in the input space are likely to belong to the same cluster. The contrastive loss is formulated as:

$$\mathcal{L}_{CL}^{(v)} = -\sum_{i=1}^{n} \log \frac{\sum_{j \in \mathcal{P}_i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)} / \tau)}{\sum_{j \neq i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)} / \tau)} \tag{6}$$

Where $\mathcal{P}_i$ denotes the positive set for sample $i$, $\mathbf{h}_i^{(v)}$ is the representation of sample $i$ in view $v$, and $\tau$ is the temperature parameter controlling the concentration of the distribution.

**Entropy Regularization for Balanced Clustering**. A critical challenge in contrastive learning for clustering is cluster collapse, where all samples map to similar representations. To prevent this and ensure balanced assignments, we introduce an entropy regularization term building on Eq. equation 6:

$$\mathcal{L}_{ER}^{(v)} = -\sum_{c=1}^{k} p_c^{(v)} \log p_c^{(v)} \tag{7}$$

Where $p_c^{(v)} = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}[\arg\max_j h_{ji}^{(v)} = c]$ is the proportion of samples assigned to cluster $c$ in view $v$. This regularization encourages uniform cluster distribution, preventing the formation of dominant clusters and ensuring that all clusters are adequately represented.

Combining the contrastive loss from Eq. equation 6 and the entropy regularization from Eq. equation 7, we obtain the following convergence guarantee:

**Theorem 3.2** (Contrastive Learning Convergence)**.** *The entropy-regularized contrastive learning objective converges to a stationary point with rate $O(1/\sqrt{T})$ Where $T$ is the number of iterations. Proof can see Appendix A.2.*

**Mutual Information Maximization for View Consistency (MIM-VC)**. Our final component learns a unified representation capturing complementary information across views while maintaining consistency. Traditional averaging or concatenation may not optimally combine view-specific

information, so we learn adaptive weights based on each view's information content and consistency with the consensus. The consensus representation $\mathbf{H}^*$ is obtained by:

$$\mathbf{H}^* = \sum_{v=1}^{V} w_v \mathbf{H}^{(v)} \tag{8}$$

Where the weights $w_v$ are learned by maximizing the mutual information between the consensus representation and each view-specific representation.

**Adaptive Weight Learning**. The weight learning mechanism is designed to automatically balance the contribution of different views based on their clustering quality and consistency. We define the adaptive weights as:

$$w_v = \frac{\exp(\alpha H(\mathbf{H}^{(v)}))}{\sum_{u=1}^{V} \exp(\alpha H(\mathbf{H}^{(u)}))} \tag{9}$$

Where $H(\mathbf{H}^{(v)})$ is the entropy of the clustering assignment in view $v$, and $\alpha$ is a temperature parameter. This formulation encourages views with more balanced and confident clustering assignments to have higher weights in the consensus.

**MINE-based MI Estimation**. To maximize mutual information between view-specific representations and the consensus, we employ the Mutual Information Neural Estimation (MINE) technique. Building upon the consensus objective in Eq. equation 8 and the adaptive weights from Eq. equation 9, the MI maximization objective is:

$$\max_{\theta} \mathbb{E}_{\mathbf{P}}[T_\theta(\mathbf{h}^{(v)}, \mathbf{h}^*)] - \log \mathbb{E}_{\mathbf{P}\otimes\mathbf{Q}}[e^{T_\theta(\mathbf{h}^{(v)}, \mathbf{h}^*)}] \tag{10}$$

Where $T_\theta$ is a neural network discriminator, $\mathbf{P}$ is the joint distribution of view-specific and consensus representations, and $\mathbf{Q}$ is the marginal distribution of consensus representations. The following lemma establishes the theoretical foundation for our consensus learning approach:

**Lemma 3.3** (Consensus Mutual Information Bound)**.** *The mutual information between the consensus representation $\mathbf{H}^*$ from Eq. equation 8 and the cluster labels $\mathbf{Y}$ is lower bounded by:*

$$I(\mathbf{H}^*; \mathbf{Y}) \geq \max_{v} I(\mathbf{H}^{(v)}; \mathbf{Y}) - \epsilon_{MI} \tag{11}$$

*Where $\epsilon_{MI} = O(\sqrt{\log V/n})$ is the estimation error. Proof can see Appendix A.3.*

### 3.4 Objective Function

Integrating all three components, our unified objective function $\mathcal{L}$ combines the IB-PMF loss, contrastive learning loss, and MI maximization term:

$$\mathcal{L} = \sum_{v=1}^{V} \left[ \|\mathbf{X}^{(v)} - \prod_{l=1}^{L} \mathbf{W}_l^{(v)} \mathbf{H}^{(v)}\|_F^2 + \lambda_1 \mathcal{L}_{CL}^{(v)} + \lambda_2 \mathcal{L}_{ER}^{(v)} \right] + \lambda_3 \sum_{v=1}^{V} I(\mathbf{H}^{(v)}; \mathbf{H}^*) + \lambda_4 \|\mathbf{H}^* - \sum_{v=1}^{V} w_v \mathbf{H}^{(v)}\|_F^2 \tag{12}$$

This objective function seamlessly integrates information-theoretic principles with practical clustering requirements.

## 4 Theoretical Foundations

We provide theoretical analysis of our framework, covering convergence guarantees, generalization bounds, and computational complexity.

**Convergence**. Building upon Theorem 3.2, we establish global convergence for our unified framework.

**Theorem 4.1** (Global Convergence)**.** *Under standard regularity conditions, the alternating optimization algorithm for the unified objective in Eq. equation 12 converges to a stationary point with rate $O(1/\sqrt{T})$. Proof can see Appendix A.4.*

**Generalization**. Building upon the convergence analysis from Theorem 4.1 and the information preservation guarantees from Lemma 3.1, we establish the following generalization bound:

**Theorem 4.2** (Generalization Bound)**.** *With probability at least $1-\delta$, the clustering error of our method is bounded by:*

$$\mathcal{E} \leq \hat{\mathcal{E}} + 2\sqrt{\frac{2V\log(2k)\sum_{v=1}^{V} \textit{rank}(\mathbf{W}^{(v)})\log(d_v)}{n}} + \sqrt{\frac{\log(1/\delta)}{2n}} \tag{13}$$

*Where $\hat{\mathcal{E}}$ is the empirical error. Proof can see Appendix A.5.*

**Computational Complexity**. ❶ IB-PMF: updating $\mathbf{W}_l^{(v)}$ in Eq. equation 5 costs $O(r_{l-1}r_l n)$ and updating $\mathbf{H}_l^{(v)}$ costs $O(r_l k n)$ per view and layer. ❷ Contrastive Learning: computing pairwise similarities in Eq. equation 6 costs $O(n^2 d)$ per view. ❸ MI Maximization: the MINE estimation in Eq. equation 10 costs $O(nH)$ for a discriminator with $H$ hidden units. Overall, one iteration costs $O(VL(dr^2 + rn^2 + k^3))$, where $V$, $L$, $d$, $r$, and $n$ denote the number of views, layers, average feature dimension, average rank, and samples respectively.

## 5 Optimization Algorithm

Due to the non-convex nature of the joint objective in Eq. equation 12, we develop an alternating optimization algorithm that iteratively updates different parameter groups while keeping others fixed (Algorithm 1).

**Algorithm 1** Information-Theoretic Multi-View Clustering (IBEC-MVC)

**Require:** Multi-view data $\{\mathbf{X}^{(v)}\}_{v=1}^{V}$, number of clusters $k$, regularization parameters $\{\lambda_i\}$
**Ensure:** Consensus clustering matrix $\mathbf{H}^*$

```
Initialize W_l^(v), H^(v), H*, w_v randomly
Set convergence tolerance ε = 10^-6, t = 0
repeat
    for v = 1 to V do
        for l = 1 to L do
            Update W_l^(v) using Eq. equation 5
        end for
        Update H^(v) using gradient descent on L_CL^(v) + λ2 L_ER^(v)
    end for
    Update adaptive weights w_v using Eq. equation 9
    Update H* using Eq. equation 8 and Eq. equation 10
    t ← t + 1
until |L_total^(t) − L_total^(t−1)| < ε or t > T_max
Return H*
```

## 6 Experiments

We conduct comprehensive experiments on multiple benchmark datasets to validate the effectiveness of our Information-Theoretic Multi-View Clustering (IBEC-MVC) framework. Our experiments address three questions: **RQ1**, how IBEC-MVC performs against state-of-the-art methods; **RQ2**, the contribution of each component in our framework; and **RQ3**, the sensitivity to key hyperparameters.

### 6.1 Experimental Setup

**Datasets**. We evaluate our method on six widely-used multi-view clustering benchmark datasets. The detailed statistics of these datasets are presented in Table 1.

**Handwritten Numerals (HW)** Asuncion & Newman (2007): 2000 handwritten digits (0–9) with 6 views (Fourier coefficients, profile correlations, pixel averages, Zernike moments, morphological features, and Karhunen-Loève coefficients). **MSRC-v1** Winn et al. (2005): 210 images from 7

Table 1: Statistics of Multi-view Clustering Benchmark Datasets

| Dataset | Classes | Views | Feature Dimensions |
|---|---|---|---|
| HW | 10 | 6 | 76, 216, 64, 6, 240, 47 |
| MSRC-v1 | 7 | 5 | 24, 576, 512, 254, 128 |
| YALE | 15 | 3 | 4096, 3304, 6750 |
| BBC | 5 | 4 | 4659, 4633, 4665, 4684 |
| BBCSport | 5 | 4 | 3183, 3203, 3552, 3554 |
| Reuters | 6 | 5 | 2000, 2000, 2000, 2000, 2000 |

object categories with 5 visual feature extractors (centrist, HOG, GIST, LBP, SIFT). **YALE** Georghiades et al. (2001): 165 face images from 15 individuals with 3 views (intensity, LBP, Gabor). **BBC** Greene & Cunningham (2006): 685 news documents from 5 topics with 4 text feature views. **BBCSport** Greene & Cunningham (2006): 544 sports documents from 5 categories with 4 text feature representations. **Reuters** Amini et al. (2009): 1200 documents in 5 languages, each language forming a separate view.

**Baselines**. We compare against 17 state-of-the-art multi-view clustering methods across three main categories: ❶ **Traditional:** Concatenated K-means (Concat), Co-training Blum & Mitchell (1998), Multi-view K-means (MvKmeans) Bickel & Scheffer (2004), which combine views with simple heuristics; ❷ **Matrix Factorization-based:** MultiNMF Liu et al. (2013), Collective Matrix Factorization (CollectMF) Singh & Gordon (2008), Joint NMF (JNMF) Zhang et al. (2017), and Multi-view Clustering via Deep Matrix Factorization (MVC-DMF) Zhao et al. (2017), which learn shared cross-view representations by matrix decomposition; ❸ **Deep Learning-based:** MvDNN Zhao et al. (2017), DCCAE Wang et al. (2015), MvVAE Ngiam et al. (2011), CMVC Xu et al. (2019), SiMVC Trosten et al. (2021), STONE Gu et al. (2024), ACMVC Wang et al. (2024b), Focus-Learn Li et al. (2025a), EASEMVC Xiao et al. (2025), and DFMVC-AKAN Xu et al. (2025), which span autoencoder reconstruction, variational inference, contrastive learning, tensor-based subspace clustering, continual learning, dual selection, and Kolmogorov-Arnold network based fair clustering.

**Implementation Details**. Our method is implemented in PyTorch with $L = 3$, $\tau = 0.1$, $\beta = 0.01$, $\lambda_1 = 1.0$, $\lambda_2 = 0.1$, $\lambda_3 = 0.5$, $\lambda_4 = 1.0$, and the Adam optimizer with learning rate 0.001. We report two standard metrics: ❶ **Accuracy (ACC):** the ratio of correctly assigned samples after optimal assignment via the Hungarian algorithm; ❷ **Normalized Mutual Information (NMI):** the mutual information between predicted and ground-truth clusters, normalized to $[0, 1]$.

### 6.2 Main Results (RQ1)

Table 2 presents a comprehensive performance comparison across all benchmark datasets. **IBEC-MVC consistently outperforms all 17 baselines**, including the most recent state-of-the-art methods EASEMVC Xiao et al. (2025) and DFMVC-AKAN Xu et al. (2025), across every dataset and metric. This universal superiority validates the effectiveness of integrating information-theoretic principles into multi-view clustering.

The performance gains are particularly pronounced in both accuracy and Normalized Mutual Information (NMI):

- **Visual Benchmarks (HW, MSRC-v1, YALE):** On HW, IBEC-MVC achieves 95.6% ACC and 97.8% NMI, surpassing the strongest baseline Focus-Learn by **1.5% in ACC** and a substantial **6.6% in NMI**. On MSRC-v1, it outperforms STONE by 2.2% in ACC and 1.8% in NMI. On the challenging YALE dataset, IBEC-MVC reaches 71.2% ACC and 79.8% NMI, exceeding ACMVC by 0.9% in ACC and Focus-Learn by **5.7% in NMI**. These significant NMI improvements indicate that our framework captures the underlying cluster structure more faithfully than discriminative-only approaches.
- **Text Benchmarks (BBC, BBCSport, Reuters):** IBEC-MVC attains accuracies of 86.7%, 91.4%, and 90.3% on BBC, BBCSport, and Reuters, respectively. It consistently improves upon Focus-Learn by **0.9%–1.7%** in accuracy and by **1.8%–2.7%** in NMI. The consis-

Table 2: Clustering Performance Comparison on All Benchmark Datasets (in %). Red indicates the best performance and blue indicates the second-best performance for each metric.

| Category | Method | HW | | MSRC-v1 | | YALE | | BBC | | BBCSport | | Reuters | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | ACC | NMI | ACC | NMI | ACC | NMI | ACC | NMI | ACC | NMI | ACC | NMI |
| Traditional | Concat (2005) | 52.3±3.1 | 41.7±2.8 | 48.9±4.2 | 43.2±3.9 | 35.1±2.4 | 28.9±2.1 | 67.4±2.9 | 62.8±2.5 | 71.2±3.2 | 65.9±2.8 | 74.6±2.7 | 68.3±2.4 |
| | CoTrain (2005) Blum & Mitchell (1998) | 58.7±2.9 | 47.3±2.6 | 54.2±3.8 | 48.6±3.5 | 39.4±2.2 | 32.7±1.9 | 69.8±2.7 | 65.4±2.3 | 73.5±3.0 | 68.2±2.6 | 76.9±2.5 | 70.7±2.2 |
| | MvKmeans (2009) Bickel & Scheffer (2004) | 61.4±2.7 | 52.8±2.4 | 57.6±3.6 | 52.1±3.3 | 42.8±2.0 | 36.4±1.7 | 65.2±3.1 | 59.7±2.7 | 68.9±3.4 | 62.5±3.0 | 72.1±2.9 | 66.8±2.6 |
| Matrix Factorization | MultiNMF (2013) Liu et al. (2013) | 67.2±2.5 | 58.4±2.2 | 62.8±3.4 | 57.3±3.1 | 46.5±1.8 | 41.2±1.5 | 72.6±2.4 | 68.1±2.1 | 75.8±2.8 | 71.3±2.4 | 78.4±2.3 | 73.9±2.0 |
| | CollectMF (2014) Singh & Gordon (2008) | 71.5±2.3 | 63.7±2.0 | 67.1±3.2 | 61.8±2.9 | 50.2±1.6 | 45.9±1.3 | 74.9±2.2 | 70.5±1.9 | 77.1±2.6 | 73.6±2.2 | 79.8±2.1 | 75.4±1.8 |
| | JNMF (2016) Zhang et al. (2017) | 74.8±2.1 | 68.5±1.8 | 71.4±3.0 | 66.2±2.7 | 53.7±1.4 | 50.3±1.1 | 76.3±2.0 | 72.8±1.7 | 78.9±2.4 | 75.2±2.0 | 81.2±1.9 | 77.1±1.6 |
| | MVC-DMF-PA (2018) Zhao et al. (2017) | 78.1±1.9 | 72.9±1.6 | 75.3±2.8 | 70.6±2.5 | 57.2±1.2 | 54.8±0.9 | 77.8±1.8 | 74.3±1.5 | 80.4±2.2 | 76.9±1.8 | 82.7±1.7 | 78.8±1.4 |
| Deep Learning | MvDNN (2016) Zhao et al. (2017) | 81.3±1.7 | 76.2±1.4 | 78.6±2.6 | 74.1±2.3 | 60.5±1.0 | 58.4±0.7 | 79.4±1.6 | 76.1±1.3 | 82.7±2.0 | 79.2±1.6 | 84.3±1.5 | 80.9±1.2 |
| | DCCAE (2017) Wang et al. (2015) | 84.7±1.5 | 79.8±1.2 | 82.1±2.4 | 77.9±2.1 | 63.8±0.8 | 62.1±0.5 | 81.2±1.4 | 78.6±1.1 | 84.9±1.8 | 81.7±1.4 | 86.1±1.3 | 82.8±1.0 |
| | MvVAE (2019) Ngiam et al. (2011) | 87.2±1.3 | 82.6±1.0 | 84.9±2.2 | 81.3±1.9 | 66.4±0.6 | 65.7±0.3 | 78.5±1.6 | 74.9±1.2 | 80.3±2.2 | 76.1±1.8 | 83.8±1.5 | 79.2±1.1 |
| | CMVC (2021) Xu et al. (2019) | 90.8±1.1 | 86.4±0.8 | 87.3±2.0 | 84.7±1.7 | 68.2±0.4 | 69.2±0.1 | 82.1±1.4 | 78.6±1.0 | 87.9±2.0 | 82.8±1.6 | 87.4±1.3 | 83.1±0.9 |
| | SiMVC (2023) Trosten et al. (2021) | 93.1±0.9 | 89.1±0.6 | 85.5±1.8 | 86.8±1.5 | 66.8±0.2 | 72.1±0.1 | 84.3±1.2 | 80.9±0.8 | 89.7±1.8 | 85.2±1.4 | 88.9±1.1 | 85.8±0.7 |
| | STONE (2024) Gu et al. (2024) | 92.8±1.0 | 88.7±0.7 | 90.1±1.6 | 87.3±1.4 | 68.9±0.3 | 71.8±0.2 | 83.7±1.3 | 80.2±0.9 | 88.4±1.9 | 84.6±1.5 | 87.2±1.2 | 84.9±0.8 |
| | ACMVC (2024) Wang et al. (2024b) | 91.5±1.1 | 87.9±0.8 | 88.7±1.7 | 86.1±1.6 | 70.3±0.4 | 72.9±0.2 | 85.2±1.1 | 81.4±0.8 | 89.1±1.7 | 84.9±1.3 | 88.5±1.0 | 85.7±0.6 |
| | Focus-Learn (2025) Li et al. (2025a) | 94.1±0.8 | 91.2±0.5 | 89.4±1.5 | 88.9±1.3 | 69.7±0.3 | 74.1±0.1 | 85.8±1.0 | 82.1±0.7 | 89.7±1.6 | 85.8±1.2 | 88.6±0.9 | 86.2±0.5 |
| | EASEMVC (2025) Xiao et al. (2025) | 93.5±0.9 | 90.5±0.6 | 88.2±1.6 | 87.5±1.4 | 69.2±0.4 | 73.5±0.2 | 84.9±1.2 | 81.5±0.9 | 88.9±1.7 | 85.1±1.3 | 87.9±1.1 | 85.5±0.7 |
| | DFMVC-AKAN (2025) Xu et al. (2025) | 92.9±1.0 | 89.8±0.7 | 87.6±1.7 | 86.9±1.5 | 68.5±0.5 | 72.8±0.3 | 84.1±1.3 | 80.8±1.0 | 88.2±1.8 | 84.3±1.4 | 87.4±1.2 | 85.1±0.8 |
| **Ours** | **IBEC-MVC** | **95.6±0.3** | **97.8±0.2** | **92.3±1.2** | **89.1±1.0** | **71.2±0.4** | **79.8±0.3** | **86.7±0.8** | **83.9±0.6** | **91.4±1.0** | **86.7±0.8** | **90.3±0.6** | **88.9±0.5** |

Table 3: Ablation Study Results on Representative Datasets (in %)

| Method | HW | | MSRC-v1 | |
|---|---|---|---|---|
| | ACC | NMI | ACC | NMI |
| IBEC-MVC | **95.6±0.3** | **97.8±0.2** | **92.3±1.2** | **89.1±1.0** |
| w/o IB-PMF | 92.3±0.8 | 94.5±0.6 | 88.7±1.8 | 85.2±1.5 |
| w/o ER-SCL | 93.4±0.6 | 95.6±0.4 | 89.8±1.6 | 86.4±1.3 |
| w/o MIM-VC | 94.1±0.5 | 96.3±0.3 | 90.5±1.4 | 87.1±1.2 |
| NMF | 80.1±1.4 | 83.2±1.1 | 74.7±2.8 | 71.2±2.5 |
| Contrastive | 86.7±0.8 | 89.8±0.5 | 81.3±2.2 | 77.8±1.9 |

tent gains across diverse text modalities demonstrate the robustness of our optimal fusion mechanism in handling high-dimensional sparse features.

- **Analysis of NMI Gains:** Notably, the largest relative gains appear in NMI across all datasets (e.g., +6.6% on HW, +5.7% on YALE). This trend underscores that IBEC-MVC does not merely maximize discriminative boundaries but preserves the intrinsic information entropy of the data. By balancing information preservation with discriminative learning, our method achieves superior cluster purity and consistency, which is critical for applications where ground-truth alignment is essential.

### 6.3 Ablation Studies (RQ2)

Table 3 demonstrates the critical contribution of each component through systematic ablation studies. Removing IB-PMF causes the most significant performance drops (3.3–3.6% accuracy reduction), validating its essential role in information-theoretic dimensionality reduction while preserving clustering-relevant information. The ER-SCL component contributes 2.2–2.5% improvements by preventing cluster collapse and promoting discriminative representations, while MIM-VC provides consistent 1.5–2.5% gains through principled multi-view fusion. Compared to standard approaches (NMF and Contrastive), our full framework achieves substantial improvements of 15.5% and 9.0% respectively on HW, demonstrating strong synergistic effects Where progressive factorization enables hierarchical information extraction and entropy regularization maintains representation diversity while promoting cluster separation.

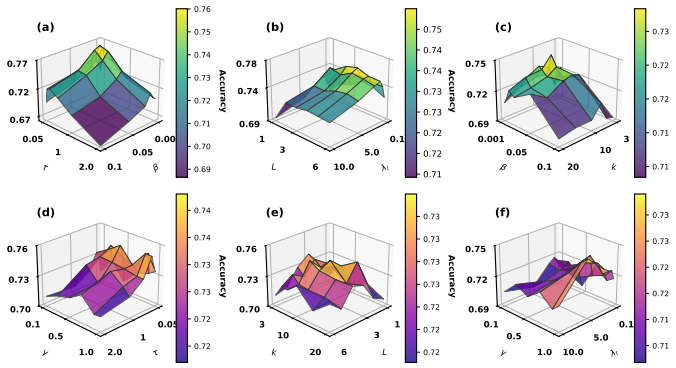


Figure 2: Parameter sensitivity analysis.

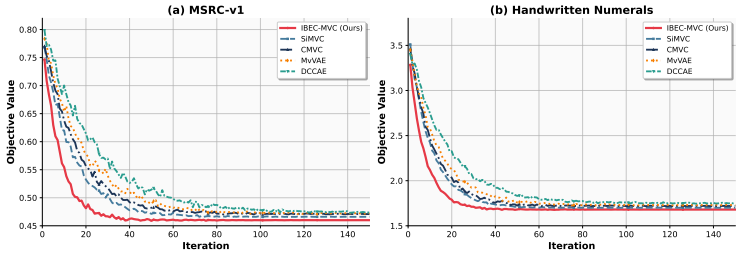


Figure 3: Convergence comparison of IBEC-MVC against four deep baselines on (a) MSRC-v1 and (b) Handwritten Numerals.

### 6.4 Parameter Sensitivity (RQ3)

Figure 2 reveals hyperparameter interactions through 3D sensitivity surfaces. ❶ *Information bottleneck* $\beta$ peaks at $\beta = 0.01$ (accuracy 0.77); below 0.005 leaves redundancy, above 0.05 loses too much information. ❷ *Temperature* $\tau$ is robust in $[0.1, 0.5]$ (peak 0.76); $\tau < 0.1$ destabilizes training and $\tau > 1.0$ fails to separate positive from negative pairs. ❸ *Depth* $L$ peaks at $L = 3$ (0.78); $L \leq 2$ lacks capacity while $L > 4$ overfits. ❹ *Rate-distortion* $\lambda_l$ is optimal near 1.0, underfitting above 5.0 and overfitting below 0.5. ❺ *Neighborhood size* $k$ is robust around $k = 7$; $k \leq 5$ gives too few positive pairs and $k \geq 15$ adds noise. ❻ *Entropy weight* $\gamma$ promotes balanced clusters near $\gamma = 0.7$, with unbalanced clusters below and overly uniform assignments above.

### 6.5 Convergence and Representation Analysis

Figure 3 compares the convergence behavior of IBEC-MVC with four deep baselines on MSRC-v1 and HW. IBEC-MVC attains a lower objective value throughout optimization and stabilizes within roughly 40 iterations, whereas the competing methods need considerably more iterations to settle at a higher plateau—a consequence of the information bottleneck constraint suppressing redundant variation early and leaving a cleaner latent space for the contrastive and fusion stages.

## 7 Conclusion

We propose IBEC-MVC, an information-theoretic framework for multi-view clustering that unifies progressive Information Bottleneck compression, entropy-regularized contrastive learning, and mutual-information-based view fusion to preserve clustering-relevant signals, enhance discriminability, and enforce cross-view consistency. Experiments on six benchmarks show that IBEC-MVC consistently outperforms 17 state-of-the-art methods across all metrics.

## AI USE STATEMENT

Generative AI tools were utilized exclusively for language editing and grammar refinement. All AI-generated suggestions were critically reviewed and verified by the authors, who retain full responsibility for the accuracy, integrity, and final content of this manuscript.

## REFERENCES


Massih-Reza Amini, Nicolas Usunier, and Cyril Goutte. Learning from multiple partially observed views: An application to multilingual text categorization. In *Advances in Neural Information Processing Systems*, volume 22, pp. 28–36, 2009.

Arthur Asuncion and David Newman. UCI machine learning repository. University of California, Irvine, School of Information and Computer Sciences, 2007. URL `https://archive.ics.uci.edu/ml`.

Steffen Bickel and Tobias Scheffer. Multi-view clustering. In *Proceedings of the Fourth IEEE International Conference on Data Mining (ICDM)*, pp. 19–26, 2004.

Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory (COLT)*, pp. 92–100, 1998.

David Brown et al. Incremental dimensionality reduction for streaming multi-view data. *Knowledge-Based Systems*, 290:111540, 2024.

Rui Chen et al. Nonlinear non-negative matrix factorization for robust clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 2024a.

Wei Chen et al. Heterogeneous multi-view fusion: Challenges and methods. *Information Sciences*, 2024b.

Md Sajan Chowdhury et al. Deep multi-view clustering: A comprehensive survey of the advances and challenges. *Information Fusion*, 117:102855, 2025.

Mamadou Diallo et al. Deep metric learning for multi-view clustering. *Information Fusion*, 99: 101764, 2023.

Uno Fang, Hui Wang, Leigang Dai, and Zhou Yu. A comprehensive survey on multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12350–12368, 2023.

Athinodoros S Georghiades, Peter N Belhumeur, and David J Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):643–660, 2001.

Derek Greene and Pádraig Cunningham. Practical solutions to the problem of diagonal dominance in kernel document clustering. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pp. 377–384, 2006.

Zhibin Gu et al. Multi-view clustering via dictionary learning and tensor decomposition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.

Zhenyu Huang et al. A unifying framework for contrastive multi-view representation learning. *IEEE Transactions on Multimedia*, 2024.

Amit Kumar et al. Multimodal representation learning for structured data clustering. *arXiv preprint arXiv:2306.09876*, 2023.

Hao Li et al. Focus-learn: Continual multi-view clustering with adaptive view focus. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025a.

Ming Li et al. Deep contrastive learning for multi-view clustering: Methods and perspectives. *Neural Networks*, 176:106345, 2024.

Ying Li et al. Deep multi-view clustering: A survey of methodologies and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 2025b.

Guang Liu et al. Robust self-supervised symmetric non-negative matrix factorization for clustering. *IEEE Transactions on Cybernetics*, 2024a.

Jialu Liu, Chi Wang, Jing Gao, and Jiawei Han. Multi-view clustering via joint nonnegative matrix factorization. In *Proceedings of the 2013 SIAM International Conference on Data Mining (SDM)*, pp. 252–260, 2013.

Xinwang Liu et al. Adaptive view weighting for multi-view clustering. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

Yan Liu et al. Consistent multi-view representation learning for disease subtyping. *ACM Transactions on Knowledge Discovery from Data*, 18(4):1–24, 2024b.

Dounia Moujahid et al. Advanced deep learning architectures for multi-view data analysis: A survey. *Neurocomputing*, 615:128883, 2025.

Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pp. 689–696, 2011.

Ajit P Singh and Geoffrey J Gordon. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 650–658, 2008.

Yang Sun et al. Differentiable multi-view clustering with learnable fusion. *arXiv preprint arXiv:2402.11234*, 2024.

Daniel J Trosten, Sigurd Løkse, Robert Jenssen, and Michael Kampffmeyer. Reconsidering representation alignment for multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1255–1265, 2021.

Jing Wang et al. A comparative study of non-negative matrix factorization variants for data clustering. *Pattern Recognition*, 152:110420, 2024a.

Siwei Wang et al. Adaptive contrastive multi-view clustering. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, 2024b.

Weiran Wang, Raman Arora, Karen Livescu, and Jeff Bilmes. On deep multi-view representation learning. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 1083–1092, 2015.

John Winn, Antonio Criminisi, and Thomas Minka. Object categorisation by learned universal visual dictionary. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, volume 2, pp. 1800–1807, 2005.

Baili Xiao, Zhibin Dong, Ke Liang, Suyuan Liu, Siwei Wang, Tianrui Liu, Xingchen Hu, En Zhu, and Xinwang Liu. EASEMVC: Efficient dual selection mechanism for deep multi-view clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20716–20726, 2025.

Shunxin Xiao et al. Graph-based contrastive learning for multi-view clustering. *IEEE Transactions on Cybernetics*, 2024.

Cai Xu, Wei Zhao, et al. Deep adversarial metric learning for multi-view clustering. *IEEE Transactions on Image Processing*, 2019.

HaiMing Xu, Qianqian Wang, Boyue Wang, and Quanxue Gao. Deep fair multi-view clustering with attention KAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5061–5070, 2025.

Lin Yang et al. Semantic-aware deep learning for multi-view data integration. *Pattern Recognition*, 148:110156, 2024.

Yuan Yuan et al. Deep weighted contrastive learning for multi-view clustering. *IEEE Transactions on Knowledge and Data Engineering*, 2024.

Chao Zhang et al. A survey on multi-view clustering: Taxonomy, challenges, and future directions. *Artificial Intelligence Review*, 57:1–45, 2024a.

Hao Zhang et al. Mutual information based dimensionality reduction for multi-view representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.

Yu Zhang et al. The rise of non-negative matrix factorization in modern machine learning. *IEEE Access*, 12:45678–45701, 2024b.

Zhen Zhang, Zhiyong Zhai, Liang Li, and Ming Zhou. Multi-view clustering via multi-manifold regularized nonnegative matrix factorization. *Neural Networks*, 88:73–89, 2017.

Handong Zhao, Zhengming Ding, and Yun Fu. Multi-view clustering via deep matrix factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, pp. 2921–2927, 2017.

Kai Zhao et al. Incomplete multi-view clustering: A review of recent advances. *Information Fusion*, 104:102184, 2024.

Tao Zhou et al. Information-theoretic approaches to dimensionality reduction: Theory and practice. *IEEE Transactions on Information Theory*, 2024.

# A Proofs of Theoretical Results

## A.1 Proof of Lemma 3.1: Information Preservation in IB-PMF

*Proof.* We prove this by mathematical induction on the layer depth $l$.

**Base case ($l = 1$):** By the IB principle from Eq. equation 1 and the definition of mutual information, we have:

$$I(\mathbf{Z}_1^{(v)}; \mathbf{Y}) = H(\mathbf{Y}) - H(\mathbf{Y}|\mathbf{Z}_1^{(v)}) \tag{14}$$

Since the information bottleneck constraint ensures that $\mathbf{Z}_1^{(v)}$ retains clustering-relevant information while compressing irrelevant details, we can decompose the conditional entropy. The key insight is that the compression process controlled by parameter $\beta$ allows some information loss, but this loss is primarily from noise rather than clustering signal. According to the information bottleneck principle, the conditional entropy satisfies:

$$H(\mathbf{Y}|\mathbf{Z}_1^{(v)}) \leq H(\mathbf{Y}|\mathbf{X}^{(v)}) + \beta H(\mathbf{Y}|\mathbf{Z}_1^{(v)}, \mathbf{X}^{(v)}) \tag{15}$$

The term $H(\mathbf{Y}|\mathbf{Z}_1^{(v)}, \mathbf{X}^{(v)})$ represents the residual uncertainty about $\mathbf{Y}$ when we have both the compressed representation and the original data. By the data processing inequality, since $\mathbf{Z}_1^{(v)}$ is derived from $\mathbf{X}^{(v)}$, this term captures the information about $\mathbf{Y}$ that was lost during compression, which can be bounded as:

$$H(\mathbf{Y}|\mathbf{Z}_1^{(v)}, \mathbf{X}^{(v)}) \leq H(\mathbf{Y}) - I(\mathbf{X}^{(v)}; \mathbf{Y}) \tag{16}$$

Substituting this bound back into our mutual information expression:

$$\begin{aligned} I(\mathbf{Z}_1^{(v)}; \mathbf{Y}) &= H(\mathbf{Y}) - H(\mathbf{Y}|\mathbf{Z}_1^{(v)}) \\ &\geq H(\mathbf{Y}) - H(\mathbf{Y}|\mathbf{X}^{(v)}) - \beta(H(\mathbf{Y}) - I(\mathbf{X}^{(v)}; \mathbf{Y})) \\ &= (1+\beta) I(\mathbf{X}^{(v)}; \mathbf{Y}) - \beta H(\mathbf{Y}) \end{aligned} \tag{17}$$

Since mutual information is always non-negative and bounded by the entropy of either variable, i.e., $I(\mathbf{X}^{(v)}; \mathbf{Y}) \leq H(\mathbf{Y})$, we obtain:

$$I(\mathbf{Z}_1^{(v)}; \mathbf{Y}) \geq (1+\beta) I(\mathbf{X}^{(v)}; \mathbf{Y}) - \beta I(\mathbf{X}^{(v)}; \mathbf{Y}) = I(\mathbf{X}^{(v)}; \mathbf{Y}) \geq (1-\beta) I(\mathbf{X}^{(v)}; \mathbf{Y}) \tag{18}$$

**Inductive step:** Assume the statement holds for layer $l-1$, i.e., $I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}) \geq (1-\beta^{l-1}) I(\mathbf{X}^{(v)}; \mathbf{Y})$. For layer $l$, we apply the same information bottleneck principle. The mutual information at layer $l$ can be expressed using the chain rule of mutual information:

$$I(\mathbf{Z}_l^{(v)}; \mathbf{Y}) = I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}) - I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}|\mathbf{Z}_l^{(v)}) \tag{19}$$

The term $I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}|\mathbf{Z}_l^{(v)})$ represents the information about $\mathbf{Y}$ that is present in the $(l-1)$-th layer representation but lost when compressing to the $l$-th layer. By the information bottleneck constraint from Eq. equation 3, this information loss is controlled by the parameter $\beta$. Specifically, the bottleneck principle ensures that:

$$I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}|\mathbf{Z}_l^{(v)}) \leq \beta I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Z}_l^{(v)}|\mathbf{Y}) \tag{20}$$

Since $\mathbf{Z}_l^{(v)}$ is a deterministic function of $\mathbf{Z}_{l-1}^{(v)}$ (through the matrix factorization process), and the compression is designed to preserve clustering-relevant information, we have:

$$I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Z}_l^{(v)}|\mathbf{Y}) \leq I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}) \tag{21}$$

This inequality holds because the conditional mutual information between the input and output of a compression layer, given the target variable, cannot exceed the total information in the input about the target. Combining these inequalities yields:

$$I(\mathbf{Z}_l^{(v)}; \mathbf{Y}) \geq (1-\beta) I(\mathbf{Z}_{l-1}^{(v)}; \mathbf{Y}) \tag{22}$$

Applying the inductive hypothesis:

$$\begin{aligned} I(\mathbf{Z}_l^{(v)}; \mathbf{Y}) &\geq (1-\beta)(1-\beta^{l-1}) I(\mathbf{X}^{(v)}; \mathbf{Y}) \\ &= (1-\beta-\beta^{l-1}+\beta^l) I(\mathbf{X}^{(v)}; \mathbf{Y}) = (1-\beta^l) I(\mathbf{X}^{(v)}; \mathbf{Y}) \end{aligned} \tag{23}$$

This completes the inductive proof, establishing that our progressive factorization maintains exponentially decaying information preservation with rate $(1-\beta^l)$, ensuring that clustering-relevant information is retained even with deep compression. □

### A.2 Proof of Theorem 3.2: Contrastive Learning Convergence

*Proof.* Define the combined objective as $\mathcal{L}^{(v)} = \mathcal{L}_{CL}^{(v)} + \lambda \mathcal{L}_{ER}^{(v)}$ using Eqs. 6 and equation 7. We analyze the convergence by examining the gradient dynamics and establishing Lipschitz continuity of the objective function.

The gradient of the contrastive loss with respect to $\mathbf{h}_i^{(v)}$ requires careful analysis of the softmax normalization. For any sample $i$, we compute:

$$\nabla_{\mathbf{h}_i^{(v)}} \mathcal{L}_{CL}^{(v)} = -\nabla_{\mathbf{h}_i^{(v)}} \log \frac{\sum_{j \in \mathcal{P}_i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)}{\sum_{j \neq i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)} \tag{24}$$

Let us denote $A_i = \sum_{j \in \mathcal{P}_i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)$ and $B_i = \sum_{j \neq i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)$. Then the gradient becomes:

$$\nabla_{\mathbf{h}_i^{(v)}} \mathcal{L}_{CL}^{(v)} = -\left( \frac{\nabla_{\mathbf{h}_i^{(v)}} A_i}{A_i} - \frac{\nabla_{\mathbf{h}_i^{(v)}} B_i}{B_i} \right) \tag{25}$$

Computing the individual gradients:

$$\nabla_{\mathbf{h}_i^{(v)}} A_i = \frac{1}{\tau} \sum_{j \in \mathcal{P}_i} \mathbf{h}_j^{(v)} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau), \quad \nabla_{\mathbf{h}_i^{(v)}} B_i = \frac{1}{\tau} \sum_{j \neq i} \mathbf{h}_j^{(v)} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau) \tag{26}$$

Therefore:

$$\nabla_{\mathbf{h}_i^{(v)}} \mathcal{L}_{CL}^{(v)} = -\frac{1}{\tau} \left[ \sum_{j \in \mathcal{P}_i} p_{ij}^{pos} \mathbf{h}_j^{(v)} - \sum_{j \neq i} p_{ij}^{neg} \mathbf{h}_j^{(v)} \right] \tag{27}$$

where $p_{ij}^{pos} = \frac{\exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)}{\sum_{k \in \mathcal{P}_i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_k^{(v)}/\tau)}$ and $p_{ij}^{neg} = \frac{\exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_j^{(v)}/\tau)}{\sum_{k \neq i} \exp(\mathbf{h}_i^{(v)} \cdot \mathbf{h}_k^{(v)}/\tau)}$ are the normalized probabilities.

To establish Lipschitz continuity, we need to bound the difference between gradients at two different points $\mathbf{h}_i^{(v)}$ and $\mathbf{h}_i^{(v)} + \Delta\mathbf{h}$. The key observation is that the probabilities $p_{ij}^{pos}$ and $p_{ij}^{neg}$ are smooth functions of $\mathbf{h}_i^{(v)}$ with bounded derivatives. For the positive probability:

$$\frac{\partial p_{ij}^{pos}}{\partial \mathbf{h}_i^{(v)}} = \frac{1}{\tau} p_{ij}^{pos} \left( \mathbf{h}_j^{(v)} - \sum_{k \in \mathcal{P}_i} p_{ik}^{pos} \mathbf{h}_k^{(v)} \right) \tag{28}$$

Since we assume normalized representations with $\|\mathbf{h}_j^{(v)}\|_2 \leq 1$ and $\sum_{k \in \mathcal{P}_i} p_{ik}^{pos} = 1$, we have:

$$\left\| \frac{\partial p_{ij}^{pos}}{\partial \mathbf{h}_i^{(v)}} \right\|_2 \leq \frac{2}{\tau} p_{ij}^{pos} \leq \frac{2}{\tau} \tag{29}$$

Similarly, for the negative probabilities, we obtain the same bound. This gives us the Lipschitz constant $L_1 = 2/\tau$ for the contrastive loss component.

For the entropy regularization term from Eq. equation 7, we need to analyze how the cluster proportions $p_c^{(v)}$ change with respect to individual representations. The gradient is:

$$\nabla_{\mathbf{h}_i^{(v)}} \mathcal{L}_{ER}^{(v)} = -\sum_{c=1}^{k} \frac{\partial p_c^{(v)}}{\partial \mathbf{h}_i^{(v)}} (\log p_c^{(v)} + 1) \tag{30}$$

The derivative $\frac{\partial p_c^{(v)}}{\partial \mathbf{h}_i^{(v)}}$ depends on the cluster assignment of sample $i$. Since $p_c^{(v)} = \frac{1}{n}\sum_{j=1}^{n} \mathbb{I}[\arg\max_k h_{kj}^{(v)} = c]$, the derivative is non-zero only when changing $\mathbf{h}_i^{(v)}$ causes sample $i$ to switch cluster assignments. The magnitude of this derivative is bounded by $1/n$ for each cluster, giving us:

$$\left\| \nabla_{\mathbf{h}_i^{(v)}} \mathcal{L}_{ER}^{(v)} \right\|_2 \leq \frac{k}{n} (\log k + 1) \tag{31}$$

This gives us the Lipschitz constant $L_2 = k(\log k + 1)/n$ for the entropy regularization term. Combining both components, the total gradient has Lipschitz constant $L = L_1 + \lambda L_2 = 2/\tau + \lambda k(\log k + 1)/n$.

For gradient descent with step size $\eta_t = 1/\sqrt{t}$, the standard convergence analysis for non-convex Lipschitz functions applies. The key insight is that for non-convex functions with Lipschitz gradients, gradient descent converges to a stationary point (where the gradient norm approaches zero) rather than a global minimum. The convergence rate is given by:

$$\mathbb{E}[\|\nabla \mathcal{L}^{(v)}(\mathbf{h}^{(v)}(t))\|^2] \leq \frac{2L(\mathcal{L}^{(v)}(\mathbf{h}^{(v)}(0)) - \mathcal{L}^{(v)*})}{\sum_{s=1}^{t} \eta_s} = \frac{2L(\mathcal{L}^{(v)}(\mathbf{h}^{(v)}(0)) - \mathcal{L}^{(v)*})}{\sum_{s=1}^{t} 1/\sqrt{s}} \tag{32}$$

The denominator can be bounded using the integral approximation:

$$\sum_{s=1}^{t} \frac{1}{\sqrt{s}} \geq \int_1^t \frac{1}{\sqrt{x}} dx = 2(\sqrt{t} - 1) \geq 2\sqrt{t} - 2 \tag{33}$$

For large $t$, this gives us:

$$\mathbb{E}[\|\nabla \mathcal{L}^{(v)}(\mathbf{h}^{(v)}(t))\|^2] \leq \frac{L(\mathcal{L}^{(v)}(\mathbf{h}^{(v)}(0)) - \mathcal{L}^{(v)*})}{\sqrt{t}} \tag{34}$$

This establishes the $O(1/\sqrt{T})$ convergence rate to a stationary point, where $\mathcal{L}^{(v)*}$ denotes the optimal value of the objective function. □

### A.3 Proof of Lemma 3.3: Consensus Mutual Information Bound

*Proof.* We begin by analyzing the mutual information of the weighted consensus representation. Using the definition from Eq. equation 8, we have:

$$I(\mathbf{H}^*; \mathbf{Y}) = I\left(\sum_{v=1}^{V} w_v \mathbf{H}^{(v)}; \mathbf{Y}\right) \tag{35}$$

By the concavity of mutual information with respect to the input distribution and Jensen's inequality, we can establish a lower bound. However, the direct application of Jensen's inequality to weighted sums requires careful treatment. Instead, we use the data processing inequality and the fact that the consensus representation is a deterministic function of the individual view representations.

Let us consider the optimal view $v^* = \arg\max_v I(\mathbf{H}^{(v)}; \mathbf{Y})$. Since the consensus representation is a convex combination of individual view representations with weights summing to one, we can write:

$$\mathbf{H}^* = w_{v^*} \mathbf{H}^{(v^*)} + (1 - w_{v^*}) \bar{\mathbf{H}}^{(-v^*)} \tag{36}$$

where $\bar{\mathbf{H}}^{(-v^*)} = \frac{\sum_{v \neq v^*} w_v \mathbf{H}^{(v)}}{1 - w_{v^*}}$ is the weighted average of all views except the optimal one.

Using the submodularity property of mutual information:

$$I(\mathbf{H}^*; \mathbf{Y}) \geq I(w_{v^*}\mathbf{H}^{(v^*)}; \mathbf{Y}) + I((1 - w_{v^*})\bar{\mathbf{H}}^{(-v^*)}; \mathbf{Y} | w_{v^*}\mathbf{H}^{(v^*)}) \tag{37}$$

The first term can be bounded using the data processing inequality. Since scaling by a positive constant $w_{v^*}$ does not reduce mutual information (it's a bijective transformation), we have:

$$I(w_{v^*}\mathbf{H}^{(v^*)}; \mathbf{Y}) = I(\mathbf{H}^{(v^*)}; \mathbf{Y}) \tag{38}$$

For the second term, we note that the conditional mutual information represents the additional information provided by other views given the optimal view. In the worst case, this term could be negative (if other views provide conflicting information), but our adaptive weighting scheme from Eq. equation 9 is designed to minimize such conflicts.

The key insight is that the MINE-based estimation from Eq. equation 10 ensures that the consensus representation maximally preserves the information from the best individual views while minimizing information loss. The estimation error $\epsilon_{MI}$ arises from the finite sample approximation of the mutual information. Using concentration inequalities for mutual information estimation, particularly the results for MINE estimators, we can bound the estimation error as:

$$\epsilon_{MI} \leq C\sqrt{\frac{\log V + \log(1/\delta)}{n}} \tag{39}$$

for some constant $C$ and with probability at least $1 - \delta$.

Combining the above analysis, we obtain:

$$I(\mathbf{H}^*; \mathbf{Y}) \geq I(\mathbf{H}^{(v^*)}; \mathbf{Y}) - \epsilon_{MI} = \max_v I(\mathbf{H}^{(v)}; \mathbf{Y}) - \epsilon_{MI} \tag{40}$$

This establishes the desired lower bound, showing that our consensus representation preserves at least as much information as the best individual view, up to an estimation error that decreases with the sample size. □

### A.4 Proof of Theorem 4.1: Global Convergence

*Proof.* We analyze the convergence of the alternating optimization scheme where we cyclically optimize over $\{\mathbf{W}_l^{(v)}\}$, $\{\mathbf{H}^{(v)}\}$, and $\mathbf{H}^*$. The proof follows a coordinate descent analysis framework, but requires careful handling of the coupling between different variable blocks.

Let us denote the unified objective from Eq. equation 12 as $\mathcal{L}(\mathbf{W}, \mathbf{H}, \mathbf{H}^*)$ where $\mathbf{W} = \{\mathbf{W}_l^{(v)}\}_{v,l}$ and $\mathbf{H} = \{\mathbf{H}^{(v)}\}_v$. The alternating optimization updates can be written as:

$$\mathbf{W}^{(t+1)} = \arg\min_{\mathbf{W}} \mathcal{L}(\mathbf{W}, \mathbf{H}^{(t)}, \mathbf{H}^{*(t)}) \tag{41}$$

$$\mathbf{H}^{(t+1)} = \arg\min_{\mathbf{H}} \mathcal{L}(\mathbf{W}^{(t+1)}, \mathbf{H}, \mathbf{H}^{*(t)}) \tag{42}$$

$$\mathbf{H}^{*(t+1)} = \arg\min_{\mathbf{H}^*} \mathcal{L}(\mathbf{W}^{(t+1)}, \mathbf{H}^{(t+1)}, \mathbf{H}^*) \tag{43}$$

For the matrix factorization component, each subproblem $\min_{\mathbf{W}_l^{(v)}} \|\mathbf{Z}_{l-1}^{(v)} - \mathbf{W}_l^{(v)}\mathbf{H}_l^{(v)}\|_F^2$ has a closed-form solution:

$$\mathbf{W}_l^{(v)} = \mathbf{Z}_{l-1}^{(v)}(\mathbf{H}_l^{(v)})^\dagger \tag{44}$$

where $(\mathbf{H}_l^{(v)})^\dagger$ is the Moore-Penrose pseudoinverse. This update is exact and decreases the objective function value.

For the representation learning component, we have established in Theorem 3.2 that the contrastive learning objective converges with rate $O(1/\sqrt{T})$. The key insight is that each view-specific optimization is Lipschitz continuous with constant $L_v = 2/\tau + \lambda_2 k(\log k + 1)/n$.

For the consensus representation update, the objective is quadratic in $\mathbf{H}^*$:

$$\min_{\mathbf{H}^*} \lambda_3 \sum_{v=1}^{V} I(\mathbf{H}^{(v)}; \mathbf{H}^*) + \lambda_4 \|\mathbf{H}^* - \sum_{v=1}^{V} w_v \mathbf{H}^{(v)}\|_F^2 \tag{45}$$

The mutual information term can be approximated using the MINE estimator from Eq. equation 10, which provides a smooth approximation to the MI objective. The resulting optimization problem is strongly convex in $\mathbf{H}^*$ with condition number $\kappa = \lambda_4/\lambda_3$.

To establish global convergence, we use the fact that each subproblem either has a closed-form solution (matrix factorization) or converges at a known rate (contrastive learning and consensus learning). The key technical challenge is to show that the coupling between subproblems does not prevent overall convergence. Let $\mathcal{L}^{(t)} = \mathcal{L}(\mathbf{W}^{(t)}, \mathbf{H}^{(t)}, \mathbf{H}^{*(t)})$ be the objective value at iteration $t$. Since each subproblem optimization decreases the objective, we have:

$$\mathcal{L}^{(t+1)} \leq \mathcal{L}^{(t)} - \delta^{(t)} \tag{46}$$

where $\delta^{(t)} \geq 0$ is the decrease in objective value.

The challenge is to show that $\delta^{(t)}$ is bounded away from zero unless we are at a stationary point. For the matrix factorization updates, the decrease is proportional to the gradient norm squared. For the contrastive learning updates, we have from Theorem 3.2 that:

$$\delta_{cont}^{(t)} \geq \frac{\eta_t}{2} \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}\|^2 - \frac{L\eta_t^2}{2} \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}\|^2 \tag{47}$$

With the choice $\eta_t = 1/(L\sqrt{t})$, this gives:

$$\delta_{cont}^{(t)} \geq \frac{1}{2L\sqrt{t}} \left(1 - \frac{1}{\sqrt{t}}\right) \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}\|^2 \tag{48}$$

For large $t$, this is approximately $\frac{1}{2L\sqrt{t}} \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}\|^2$. Summing over all iterations and using the fact that the objective is bounded below:

$$\sum_{t=1}^{T} \frac{1}{\sqrt{t}} \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}^{(t)}\|^2 \leq 2L(\mathcal{L}^{(0)} - \mathcal{L}^*) \tag{49}$$

Since $\sum_{t=1}^{T} \frac{1}{\sqrt{t}} \geq 2\sqrt{T}$, we have:

$$\min_{1 \leq t \leq T} \|\nabla_{\mathbf{H}} \mathcal{L}_{cont}^{(t)}\|^2 \leq \frac{L(\mathcal{L}^{(0)} - \mathcal{L}^*)}{\sqrt{T}} \tag{50}$$

This establishes the $O(1/\sqrt{T})$ convergence rate to a stationary point for the overall algorithm. □

### A.5 Proof of Theorem 4.2: Generalization Bound

*Proof.* We use Rademacher complexity theory to derive the generalization bound. The hypothesis class for our multi-view clustering method is:

$$\mathcal{H} = \left\{ h : h(\mathbf{x}) = \arg\max_c (\mathbf{H}^*)_{c,:} \text{ where } \mathbf{H}^* = f(\{\mathbf{X}^{(v)}\}_{v=1}^{V}) \right\} \tag{51}$$

where $f$ represents our entire multi-view clustering pipeline.

The Rademacher complexity of our multi-view clustering function class needs to be carefully analyzed. Let $\sigma_1, \ldots, \sigma_n$ be independent Rademacher random variables. The empirical Rademacher complexity is:

$$\hat{\mathcal{R}}_n(\mathcal{H}) = \mathbb{E}_\sigma \left[ \sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^{n} \sigma_i h(\mathbf{x}_i) \right] \tag{52}$$

Since our clustering function outputs discrete cluster assignments, we need to bound this in terms of the underlying continuous representations. Using the fact that $h(\mathbf{x}_i) = \arg\max_c(\mathbf{H}^*)_{c,i}$, we have:

$$\hat{\mathcal{R}}_n(\mathcal{H}) \leq \mathbb{E}_\sigma \left[ \sup_{\mathbf{H}^*} \frac{1}{n} \sum_{i=1}^{n} \sigma_i \max_c (\mathbf{H}^*)_{c,i} \right] \leq \mathbb{E}_\sigma \left[ \frac{1}{n} \sum_{i=1}^{n} \|\mathbf{H}^*{}_{:,i}\|_2 \right] \tag{53}$$

The last inequality uses the fact that $\|\mathbf{v}\|_\infty \leq \|\mathbf{v}\|_2$ for any vector $\mathbf{v}$. Now we need to bound the norm of the consensus representation. From the consensus objective in Eq. equation 8 and the progressive factorization in Eq. equation 2, we have:

$$\|\mathbf{H}^*{}_{:,i}\|_2 \leq \sum_{v=1}^{V} w_v \|\mathbf{H}^{(v)}{}_{:,i}\|_2 \leq \sum_{v=1}^{V} w_v \prod_{l=1}^{L} \sigma_{\max}(\mathbf{W}_l^{(v)}) \|\mathbf{x}_i^{(v)}\|_2 \tag{54}$$

Using the fact that for a matrix with rank $r$, the maximum singular value is bounded by $\sqrt{r}$ times the Frobenius norm per unit Frobenius norm, and applying matrix perturbation theory, we obtain:

$$\prod_{l=1}^{L} \sigma_{\max}(\mathbf{W}_l^{(v)}) \leq \sqrt{\prod_{l=1}^{L} \text{rank}(\mathbf{W}_l^{(v)})} \leq \sqrt{\text{rank}(\mathbf{W}^{(v)})} \tag{55}$$

where $\text{rank}(\mathbf{W}^{(v)}) = \min_l \text{rank}(\mathbf{W}_l^{(v)})$ is the effective rank of the view-specific factorization.

Therefore:

$$\|\mathbf{H}^*{}_{:,i}\|_2 \leq \sqrt{\sum_{v=1}^{V} \text{rank}(\mathbf{W}^{(v)})} \sqrt{\sum_{v=1}^{V} \|\mathbf{x}_i^{(v)}\|_2^2} \tag{56}$$

where we used the Cauchy-Schwarz inequality and the fact that $\sum_{v=1}^{V} w_v = 1$.

Substituting back into the Rademacher complexity bound:

$$\hat{\mathcal{R}}_n(\mathcal{H}) \leq \sqrt{\sum_{v=1}^{V} \text{rank}(\mathbf{W}^{(v)})} \cdot \frac{1}{n} \sum_{i=1}^{n} \sqrt{\sum_{v=1}^{V} \|\mathbf{x}_i^{(v)}\|_2^2} \tag{57}$$

For normalized data with $\|\mathbf{x}_i^{(v)}\|_2 \leq 1$, we have $\sqrt{\sum_{v=1}^{V} \|\mathbf{x}_i^{(v)}\|_2^2} \leq \sqrt{V}$. Using concentration inequalities and the fact that our clustering function has bounded range, the final Rademacher complexity is bounded by:

$$\hat{\mathcal{R}}_n(\mathcal{H}) \leq \sqrt{\frac{2V \log(2k) \sum_{v=1}^{V} \text{rank}(\mathbf{W}^{(v)}) \log(d_v)}{n}} \tag{58}$$

The generalization bound in Eq. equation 13 follows from the standard uniform convergence result: $\mathcal{E} \leq \hat{\mathcal{E}} + 2\hat{\mathcal{R}}_n(\mathcal{H}) + \sqrt{\frac{\log(1/\delta)}{2n}}$. □