Novelty Assessment Report

Paper: Calibrated Information Bottleneck for Trusted Multi-modal Clustering
PDF URL: https://openreview.net/pdf?id=iedlZOdI0d
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Information Bottleneck (IB) Theory is renowned for its ability to learn simple, compact, and effective data representations. In multi-modal clustering, IB theory effectively eliminates interfering redundancy and noise from multi-modal data, while maximally preserving the discriminative information. Existing IB-based multi-modal clustering methods suffer from low-quality pseudo-labels and over-reliance on accurate Mutual Information (MI) estimation, which is known to be challenging. Moreover, unreliable or noisy pseudo-labels may lead to an overconfident clustering outcome. To address these challenges, this paper proposes a novel CaLibrated Information Bottleneck (CLIB) framework designed to learn a clustering that is both accurate and trustworthy. We build a parallel multi-head network architecture—incorporating one primary cluster head and several modality-specific calibration heads—which achieves three key goals: namely, calibrating for the distortions introduced by biased MI estimation thus improving the stability of IB, constructing reliable target variables for IB from multiple modalities and producing a trustworthy clustering result. Notably, we design a dynamic pseudo-label selection strategy based on information redundancy theory to extract high-quality pseudo-labels, thereby enhancing training stability. Experimental results demonstrate that our model not only achieves state-of-the-art clustering accuracy on multiple benchmark datasets but also exhibits excellent performance on the expected calibration error metric.

Core Task Landscape

This paper addresses: **Calibrated Multi-Modal Clustering with Information Bottleneck Theory**

A total of **24 papers** were analyzed and organized into a taxonomy with **16 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Information Bottleneck Architectures for Multi-Modal Clustering**
- **Information Decomposition and Fusion Strategies**
- **Calibration and Reliability Enhancement Mechanisms**
- **Theoretical Foundations and Novel Information Measures**
- **Domain-Specific Applications and Cross-Task Extensions**

Complete Taxonomy Tree

- Calibrated Multi-Modal Clustering with Information Bottleneck Theory Survey Taxonomy
- Information Bottleneck Architectures for Multi-Modal Clustering
    - Dual and Multi-Path Information Bottleneck Networks (3 papers)
    - [4] Super Deep Contrastive Information Bottleneck for Multi-modal Clustering (Zhengzheng Lou, 2025) View paper
    - [11] Robust incomplete multi-modal clustering with interpolation enhancement and dual-path contrastive optimization (B Han, 2025) View paper
    - [18] DMIB: Dual-Correlated Multivariate Information Bottleneck for Multiview Clustering. (Shizhe Hu, 2022) View paper
    - Deterministic and Differentiable Information Bottleneck Formulations (2 papers)
    - [20] Information bottleneck fusion for deep multi-view clustering (Jie Hu, 2024) View paper
    - [21] Differentiable Information Bottleneck for Deterministic Multi-View Clustering (Xiaoqiang Yan, 2024) View paper
    - Hierarchical and Propagating Information Bottleneck Frameworks (1 papers)
    - [9] Multiview Clustering With Propagating Information Bottleneck (Shizhe Hu, 2023) View paper
- Information Decomposition and Fusion Strategies
    - Shared-Private Information Separation (2 papers)
    - [10] Robust Biometric Recognition via InformationBottleneck Multi-Modal Feature Fusion (Yan Shen, 2024) View paper
    - [17] Shared-private information bottleneck method for cross-modal clustering (Xiaoqiang Yan, 2019) View paper
    - Consistent-Complementary-Private Information Disentanglement (2 papers)
    - [2] Multi-aspect Self-guided Deep Information Bottleneck for Multi-modal Clustering (Shizhe Hu, 2025) View paper
    - [6] Disentangling consistent, complementary and modality-private information via twin information bottleneck for multi-modal clustering (Ming-Yuan Li, 2025) View paper
    - Cross-Modal Correlation Mining with Information Bottleneck (1 papers)
    - [3] Cross-Modal Clustering With Deep Correlated Information Bottleneck Method (Xiaoqiang Yan, 2023) View paper
- Calibration and Reliability Enhancement Mechanisms
    - Multi-Head Calibration and Pseudo-Label Selection ★ (2 papers)
    - [0] Calibrated Information Bottleneck for Trusted Multi-modal Clustering (Anon et al., 2026) View paper
    - [1] Mutual Calibration Network for Multi-view Clustering (Yuang Xiao, 2025) View paper
    - Peer-Review and Self-Supervised Calibration (2 papers)

    - [5] A Peer-review Look on Multi-modal Clustering: An Information Bottleneck Realization Method (Zhengzheng Lou, 2025) View paper
    - [8] Self-supervised Weighted Information Bottleneck for Multi-view Clustering (Pavel Naumov, 2024) View paper
    - Adversarial Robustness and Defense Mechanisms (1 papers)
    - [23] GUARD: General Unsupervised Adversarial Robust Defense for Deep Multi-View Clustering via Information Bottleneck (H Huang, n.d.) View paper
- Theoretical Foundations and Novel Information Measures
    - Wyner Common Information and Bipartite Extensions (1 papers)
    - [7] Efficient Solvers for Wyner Common Information With Application to Multi-Modal Clustering (Teng-Hui Huang, 2024) View paper
    - Parameter-Free and Weighted Information Bottleneck (1 papers)
    - [13] A Parameter-free Multi-view Information Bottleneck Clustering Method by Cross-view Weighting (Shizhe Hu, 2022) View paper
    - Unified Frameworks for Complete and Incomplete Multi-View Data (2 papers)
    - [15] Incomplete multi-view data clustering with hidden data mining and fusion techniques (Liu, 2023) View paper
    - [16] Unifying complete and incomplete multi-view clustering through an information-theoretic generative model. (Yanghang Zheng, 2024) View paper
- Domain-Specific Applications and Cross-Task Extensions
    - Cross-View Action Recognition and Video Analysis (1 papers)
    - [14] Multi-task information bottleneck co-clustering for unsupervised cross-view human action categorization (Xiaoqiang Yan, 2020) View paper
    - Multi-Modal Entity Alignment and Knowledge Graphs (1 papers)
    - [12] IBMEA: Exploring Variational Information Bottleneck for Multi-modal Entity Alignment (Taoyu Su, 2024) View paper
    - Neuropsychiatric Disorder Analysis and Biomedical Applications (1 papers)
    - [19] Multi-modal Spatial-modality Attentive Fusion for Studying Neuropsychiatric Disorders (M. A. Rahaman, 2024) View paper
    - Visual Content Structuring and Shared-Neighbours Methods (2 papers)
    - [22] Shared-Neighbours methods for visual content structuring and mining (T Hesis, 2012) View paper
    - [24] PHD THESIS Shared-Neighbours methods for visual content structuring and mining (HAMZAOUI, 2012) View paper

Narrative

Core task: calibrated multi-modal clustering with information bottleneck theory. The field organizes around five main branches that reflect distinct methodological emphases. Information Bottleneck Architectures for Multi-Modal Clustering explores how to design neural encoders that compress multi-view data while preserving cluster-relevant information, often drawing on variants of the information bottleneck principle such as Twin Information Bottleneck[6] and Peer-review Information Bottleneck[5]. Information Decomposition and Fusion Strategies examines how to separate shared versus private components across modalities, with works like Shared-private Bottleneck[17] and Wyner Common Information[7] providing theoretical grounding. Calibration and Reliability Enhancement Mechanisms focuses on improving the trustworthiness of cluster assignments through techniques such as multi-head calibration and pseudo-label selection, exemplified by Mutual Calibration Network[1]. Theoretical Foundations and Novel Information Measures develops new mathematical tools and bounds for multi-view learning, while Domain-Specific Applications and Cross-Task Extensions demonstrates how these principles transfer to biometric recognition, entity alignment, and other specialized settings.

A particularly active line of work centers on calibration strategies that refine pseudo-labels and mitigate overconfident predictions in unsupervised settings. Mutual Calibration Network[1] and Multi-aspect Self-guided[2] both address the challenge of noisy cluster assignments by leveraging cross-view agreement and self-supervision, yet they differ in whether calibration occurs through mutual correction or through aspect-specific guidance. Calibrated Information Bottleneck[0] sits naturally within this calibration-focused branch, emphasizing multi-head architectures that jointly optimize compression and reliability. Compared to Mutual Calibration Network[1], which primarily uses cross-modal consistency checks, Calibrated Information Bottleneck[0] integrates information-theoretic constraints more tightly into the calibration process itself. Meanwhile, works in the information decomposition branch, such as Shared-private Bottleneck[17], tackle orthogonal questions about how to disentangle modality-specific noise from shared semantic structure, highlighting an ongoing tension between achieving tight compression and preserving interpretable, calibrated cluster representations.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Mutual Calibration Network for Multi-view Clustering

Authors: Yuang Xiao, Chang Tang, Xiao Zheng, Weiqing Yan, Yuanyuan Liu, et al. (6 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

â¦ The integration of the information bottleneck theory and â¦ SIB-MSC [42] is the first to introduce the information bottleneck â¦ a mutual calibration network (McMVC) for multi-view clustering, â¦

Relationship Analysis

Both papers belong to the Multi-Head Calibration and Pseudo-Label Selection category, employing multiple heads and calibration strategies to improve clustering reliability. They overlap in using calibration mechanisms to address overconfidence and enhance clustering stability in multi-modal settings. However, the original paper (CLIB) focuses on calibrating Information Bottleneck theory with dynamic pseudo-label selection based on information redundancy, while the candidate paper (McMVC) emphasizes mutual calibration between views without the explicit IB framework and uses different calibration strategies for multi-view clustering.

Contributions Analysis

Overall novelty summary. The paper proposes a Calibrated Information Bottleneck (CLIB) framework that combines multi-head calibration with dynamic pseudo-label selection for multi-modal clustering. It resides in the 'Multi-Head Calibration and Pseudo-Label Selection' leaf, which contains only one sibling paper (Mutual Calibration Network). This leaf sits within the broader 'Calibration and Reliability Enhancement Mechanisms' branch, indicating a moderately sparse research direction focused specifically on improving clustering trustworthiness through architectural calibration strategies rather than general information bottleneck design.

The taxonomy reveals that calibration-focused methods occupy one of five major branches in this field. Neighboring leaves include 'Peer-Review and Self-Supervised Calibration' and 'Adversarial Robustness and Defense Mechanisms', which address reliability through different mechanisms (self-supervision vs. adversarial training). The sibling branches—'Information Bottleneck Architectures' and 'Information Decomposition and Fusion Strategies'—tackle orthogonal challenges such as dual-path network design and shared-private information separation. CLIB's emphasis on multi-head calibration distinguishes it from these architectural and decomposition-focused approaches, positioning it at the intersection of reliability enhancement and information-theoretic compression.

Among twenty-five candidates examined, the first contribution (calibrated IB with dynamic pseudo-labels) shows overlap with two prior works, while the second (MI estimation bias mitigation) and third (trustworthy clustering with low ECE) contributions examined ten candidates each with no clear refutations. The dynamic pseudo-label selection mechanism appears to have more substantial prior work in the limited search scope, particularly from the sibling Mutual Calibration Network paper. The calibration mechanism addressing MI estimation bias and the trustworthy clustering objective appear more distinctive within the examined candidate set, though the search scope remains constrained to top-K semantic matches.

Based on the limited literature search of twenty-five candidates, the work introduces a novel integration of multi-head calibration with information bottleneck principles in a relatively sparse research direction. The calibration mechanism for MI estimation bias and the trustworthy clustering formulation appear less explored in the examined candidates, while the dynamic pseudo-label selection shows more overlap with existing calibration-focused methods. The analysis reflects top-K semantic search results and does not claim exhaustive coverage of all relevant prior work in multi-modal clustering or information bottleneck theory.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Calibrated Information Bottleneck framework with dynamic pseudo-label selection

Description: The authors introduce a novel framework that applies Information Bottleneck theory with a dynamic pseudo-label selection strategy based on information redundancy. This mechanism filters high-quality pseudo-labels to provide reliable target variables for IB, thereby improving the stability and robustness of feature extraction in multi-modal clustering.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Self-supervised Weighted Information Bottleneck for Multi-view Clustering

URL: View paper

Prior Art Analysis

Self-supervised Weighted[8] demonstrates that prior work exists on applying Information Bottleneck theory with pseudo-label mechanisms for multi-view clustering. The candidate paper explicitly describes incorporating self-supervised pseudo-label information into weighted Information Bottleneck for multi-view clustering, and claims to be 'the first work incorporating the self-supervised learning into weighted multi-view clustering.' Both papers address the fundamental challenge of using pseudo-labels to guide Information Bottleneck-based clustering, though they employ different specific mechanisms (the original uses dynamic selection based on information redundancy, while the candidate integrates pseudo-labels into weight learning).

Evidence

Evidence 1 - **Rationale**: Both papers identify the use of pseudo-labels in Information Bottleneck-based clustering as a key approach. The candidate explicitly proposes incorporating pseudo-label self-supervision into weighted Information Bottleneck, demonstrating prior work on combining IB with pseudo-label mechanisms. - **Original**: in current multi-modal clustering approaches utilizing ib (hu et al., 2025; yan et al., 2024), model-generated pseudo-labels are commonly used for guidance. a critical issue arises because these pseudo-labels are often noisy, particularly during the initial training stages. - **Candidate**: most existing weighted mvcs only consider the quality of each view, ignoring the vital role of pseudo label self-supervision information in weight learning. in this work, we propose a novel self-supervised weighted information bottleneck (swib) method for solving the multi-view clustering problem.

Evidence 2 - **Rationale**: Both papers describe mechanisms for handling pseudo-labels within Information Bottleneck frameworks for clustering. The candidate's approach of incorporating self-supervised pseudo-label information into weight learning represents a prior implementation of IB with pseudo-label guidance, though using a different technical approach than the original's dynamic selection strategy. - **Original**: we design a dynamic pseudolabel selection strategy based on information redundancy theory to extract highquality pseudo-labels, thereby enhancing training stability. - **Candidate**: it combines the weighted information from different views based on information bottleneck theory, and the view weight learning mechanism is newly designed by simultaneously taking into accounting both the quality of view-contained information and the self-supervised information on the data partition...

Evidence 3 - **Rationale**: Both papers claim to be 'first' in applying Information Bottleneck with specific pseudo-label mechanisms for multi-view/multi-modal clustering. The candidate's claim of being first to incorporate self-supervised learning (pseudo-labels) into weighted multi-view clustering directly challenges the original's novelty claim about being first to apply IB with pseudo-label selection for multi-modal clustering. - **Original**: to our knowledge, this is the first work addressing the trusted multi-modal clustering problem with calibrated ib framework. - **Candidate**: to our knowledge, this is the first work incorporating the self-supervised learning into weighted multi-view clustering.

---

2. Multi-aspect Self-guided Deep Information Bottleneck for Multi-modal Clustering

URL: View paper

Prior Art Analysis

Multi-aspect Self-guided[2] demonstrates prior work that applies Information Bottleneck theory with dynamic pseudo-label selection for multi-modal clustering. Both papers use IB theory to eliminate redundant information while preserving discriminative features, and both employ dynamic mechanisms to filter high-quality pseudo-labels. Multi-aspect Self-guided[2] explicitly describes using 'self-supervised pseudo label information' as guidance and implements a quality scoring mechanism to select reliable pseudo-labels, which directly parallels the original paper's claimed novelty of dynamic pseudo-label selection based on information redundancy.

Evidence

Evidence 1 - **Rationale**: Both papers apply Information Bottleneck theory to multi-modal clustering with the goal of extracting discriminative features while eliminating redundancy. - **Original**: we design a dynamic pseudolabel selection strategy based on information redundancy theory to extract highquality pseudo-labels, thereby enhancing training stability. - **Candidate**: msdib integrates information compression and preservation to extract discriminative features from heterogeneous data through mutual guidance between feature and clustering information. the information bottleneck (hu et al. 2024a) retains features relevant to target y by creating a compressed represe...

Evidence 2 - **Rationale**: Multi-aspect Self-guided[2] explicitly uses self-supervised pseudo labels as one of three guidance aspects in their IB-based framework, demonstrating prior work on pseudo-label integration with IB theory. - **Original**: notably, we design a dynamic pseudolabel selection strategy based on information redundancy theory to extract highquality pseudo-labels, thereby enhancing training stability. - **Candidate**: msdib employs three information aspects: private representations for each modality, a shared representation across modalities, and self-supervised pseudo labels.

Evidence 3 - **Rationale**: Both papers use information-theoretic measures (entropy/mutual information) to assess quality and guide selection mechanisms in their IB frameworks. - **Original**: the sample quality score,s(p), is designed to be negatively correlated with the entropyh(p), as shown by its derivative: ∂s(p) ∂h(p) =- 1 2hmax <0 the number of high-quality samples selected for modalitym, denotedk sel m =jpn i=1 s(pcal i,m) k - **Candidate**: to achieve our goal, we constrain the compression of information between inputs and features by minimizing the mutual information. therefore, the objective function of this part is formulated by lic = mx i=1 i(xi, hi) + mx i=1 mx j=i+1 1i=ji(hi, hj)

Evidence 4 - **Rationale**: Both papers describe mechanisms for generating and utilizing pseudo-labels within their clustering frameworks, with quality assessment and iterative refinement. - **Original**: our screening mechanism is a variant of information redundancy. the quality score for a single sample is: s(p) = 1h(p) 2xh max (9) while computing quality scores, samples are ranked according to their top probability values, and the topk sel m =⌊ pn i=1 s(pi,m)⌋samples from modalitymare selected int... - **Candidate**: to obtain modality-specific cluster assignments{y i}m i=1, we employ a multi-layer perceptron with a softmax output layer, thus mapping features into a cluster space (space for clustering with pseudo labels). within this space, features are iteratively refined to promote intra-cluster compactness an...

3. Learning Compact Semantic Information for Incomplete Multi-View Missing Multi-Label Classification

URL: View paper

Brief Assessment

Compact Semantic Missing[26] focuses on incomplete multi-view missing multi-label classification using mutual information enhancement and soft pseudo-label cross-imputation. The original paper addresses multi-modal clustering with calibrated IB and dynamic pseudo-label selection based on information redundancy. These are distinct application domains with different technical approaches to pseudo-label generation.

4. Dual global information guidance for deep contrastive multi-modal clustering

URL: View paper

Brief Assessment

Dual Global Guidance[25] focuses on contrastive learning with global information guidance for multi-modal clustering, not on Information Bottleneck theory with dynamic pseudo-label selection mechanisms as proposed in the original paper.

5. Robust incomplete multi-modal clustering with interpolation enhancement and dual-path contrastive optimization

URL: View paper

Brief Assessment

Robust Incomplete Interpolation[11] focuses on incomplete multi-modal clustering with interpolation enhancement and gated residual mechanisms, not on Information Bottleneck theory with dynamic pseudo-label selection for multi-modal clustering. The technical approaches and problem formulations are fundamentally different.

Contribution 2: Calibration mechanism to mitigate MI estimation bias

Description: The authors propose a parallel multi-head architecture with modality-specific calibration heads that can correct biases in mutual information estimation by leveraging cross-modal information. This is the first work to introduce calibration for addressing performance issues in IB arising from inaccurate MI estimation.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Calibration bottleneck: Over-compressed representations are less calibratable

URL: View paper

Brief Assessment

Calibration Bottleneck[39] addresses calibration of neural network confidence outputs in supervised classification, focusing on over-compression of representations. The original paper's calibration mechanism targets mutual information estimation bias in multi-modal clustering with Information Bottleneck theory—a fundamentally different problem domain and technical approach.

2. Debiased representation learning in recommendation via information bottleneck

URL: View paper

Brief Assessment

Debiased Recommendation Bottleneck[37] addresses confounding bias in recommender systems through independent component learning, not calibration mechanisms for correcting MI estimation bias in multi-modal clustering contexts.

3. Learning Fair Graph Representations with Multi-view Information Bottleneck

URL: View paper

Brief Assessment

Fair Graph Representations[41] addresses fairness in graph neural networks through multi-view information bottleneck for bias mitigation, not calibration mechanisms for correcting MI estimation bias in clustering frameworks.

4. Loss or gain: Hierarchical conditional information bottleneck approach for incomplete time series classification

URL: View paper

Brief Assessment

Hierarchical Conditional Bottleneck[40] focuses on incomplete time series classification with task-driven imputation, not on calibrating MI estimation bias in multi-modal clustering contexts.

5. Estimating Information Flow in DNNs

URL: View paper

Brief Assessment

Information Flow DNNs[44] focuses on estimating mutual information in DNNs through a noisy framework and does not propose calibration mechanisms for correcting MI estimation bias. The paper addresses MI estimation accuracy but not calibration for bias correction in Information Bottleneck methods.

6. LM: Mutual Information Scaling Law for Long-Context Language Modeling

URL: View paper

Brief Assessment

Long-Context Scaling Law[38] addresses mutual information scaling in long-context language modeling, not calibration mechanisms for correcting MI estimation bias in Information Bottleneck clustering methods.

7. Scalable Mutual Information Estimation using Dependence Graphs

URL: View paper

Brief Assessment

Dependence Graphs Estimation[45] focuses on computational efficiency of MI estimation itself, proposing a linear-time estimator (EDGE) for general MI computation. It does not address calibration mechanisms for correcting biases in Information Bottleneck methods or multi-modal clustering contexts.

---

8. Information theoretic counterfactual learning from missing-not-at-random feedback

URL: View paper

Brief Assessment

Counterfactual Missing Feedback[43] addresses counterfactual learning in recommender systems with MNAR data using information bottleneck theory, but does not propose calibration mechanisms for correcting mutual information estimation bias. The paper focuses on contrastive information regularization between factual and counterfactual domains, not on calibrating MI estimators themselves.

---

9. DICE: Diversity in Deep Ensembles via Conditional Redundancy Adversarial Estimation

URL: View paper

Brief Assessment

DICE Conditional Redundancy[46] focuses on ensemble diversity through conditional redundancy reduction in deep ensembles, not on calibrating mutual information estimation bias in Information Bottleneck methods for multi-modal clustering.

---

10. Analysis of Information Transfer Mechanism in Knowledge Distillation from an Information Theory Perspective

URL: View paper

Brief Assessment

Knowledge Distillation Transfer[42] focuses on knowledge distillation between teacher-student models using information theory, not on calibrating mutual information estimation bias in Information Bottleneck methods for multi-modal clustering.

Contribution 3: Trustworthy clustering with low Expected Calibration Error

Description: The framework produces clustering results that are both accurate and trustworthy by reducing model overconfidence. The calibration mechanism enables the model to achieve substantially lower ECE values while maintaining high clustering accuracy, enhancing the trustworthiness of the IB framework.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Multimodal Emotion Recognition Calibration in Conversations

URL: View paper

Brief Assessment

Emotion Recognition Calibration[30] focuses on calibrating confidence in multimodal emotion recognition in conversations, not multi-modal clustering. The tasks and domains are fundamentally different.

---

2. Unveiling uncertainty: A deep dive into calibration and performance of multimodal large language models

URL: View paper

Brief Assessment

Unveiling Uncertainty Calibration[31] focuses on calibration of multimodal large language models (MLLMs) for VQA tasks, not multi-modal clustering frameworks. The candidate addresses ECE in classification/QA settings, while the original work applies calibration mechanisms to information bottleneck-based clustering.

---

3. Calibrating Uncertainty Quantification of Multi-Modal LLMs using Grounding

URL: View paper

Brief Assessment

Grounding Uncertainty Quantification[27] focuses on calibrating multi-modal LLMs for question-answering tasks using grounding and temperature scaling, not on multi-modal clustering with information bottleneck frameworks.

---

4. DealMVC: Dual Contrastive Calibration for Multi-view Clustering

URL: View paper

Brief Assessment

DealMVC Dual Contrastive[35] focuses on contrastive calibration for view consistency in multi-view clustering, not on reducing Expected Calibration Error or addressing model overconfidence as a trustworthiness metric.

---

5. Multi-Modal Learning with Bayesian-Oriented Gradient Calibration

URL: View paper

Brief Assessment

Bayesian-Oriented Gradient Calibration[36] focuses on gradient uncertainty modeling in multi-modal learning for predictive tasks, not on clustering calibration or ECE reduction in unsupervised clustering frameworks.

---

6. Pseudo-label calibration semi-supervised multi-modal entity alignment

URL: View paper

Brief Assessment

Pseudo-label Entity Alignment[28] focuses on multi-modal entity alignment in knowledge graphs using pseudo-label calibration for semi-supervised learning, not on general multi-modal clustering with ECE reduction. The tasks and methodologies are fundamentally different.

---

7. COLD fusion: Calibrated and ordinal latent distribution fusion for uncertainty-aware multimodal emotion recognition

URL: View paper

Brief Assessment

COLD Fusion[33] focuses on uncertainty-aware emotion recognition from face and voice modalities, not clustering tasks. The calibration mechanism addresses predictive uncertainty in supervised emotion recognition, which is fundamentally different from the unsupervised clustering calibration proposed in the original paper.

---

8. Does Adding a Modality Really Make Positive Impacts in Incomplete Multi-Modal Brain Tumor Segmentation?

URL: View paper

Brief Assessment

Incomplete Brain Tumor[32] focuses on multi-modal brain tumor segmentation with positive-negative impact region calibration, not on clustering tasks or Expected Calibration Error reduction in clustering frameworks.

---

9. Calibrating multimodal learning

URL: View paper

Brief Assessment

Calibrating Multimodal Learning[29] focuses on calibrating confidence estimation in multimodal classification tasks, not clustering. The paper addresses confidence ranking relationships when modalities are removed in supervised/semi-supervised settings, which is fundamentally different from the original paper's unsupervised clustering framework with pseudo-labels and information bottleneck theory.

---

10. Calibrating class weights with multi-modal information for partial video domain adaptation

URL: View paper

Brief Assessment

Class Weights Calibration[34] focuses on partial video domain adaptation with class weight calibration for cross-domain video classification, not on multi-modal clustering with ECE reduction. The technical contexts are fundamentally different.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Calibrated Information Bottleneck for Trusted Multi-modal Clustering View paper
- [1] Mutual Calibration Network for Multi-view Clustering View paper
- [2] Multi-aspect Self-guided Deep Information Bottleneck for Multi-modal Clustering View paper
- [3] Cross-Modal Clustering With Deep Correlated Information Bottleneck Method View paper
- [4] Super Deep Contrastive Information Bottleneck for Multi-modal Clustering View paper
- [5] A Peer-review Look on Multi-modal Clustering: An Information Bottleneck Realization Method View paper
- [6] Disentangling consistent, complementary and modality-private information via twin information bottleneck for multi-modal clustering View paper
- [7] Efficient Solvers for Wyner Common Information With Application to Multi-Modal Clustering View paper
- [8] Self-supervised Weighted Information Bottleneck for Multi-view Clustering View paper
- [9] Multiview Clustering With Propagating Information Bottleneck View paper
- [10] Robust Biometric Recognition via InformationBottleneck Multi-Modal Feature Fusion View paper
- [11] Robust incomplete multi-modal clustering with interpolation enhancement and dual-path contrastive optimization View paper
- [12] IBMEA: Exploring Variational Information Bottleneck for Multi-modal Entity Alignment View paper
- [13] A Parameter-free Multi-view Information Bottleneck Clustering Method by Cross-view Weighting View paper
- [14] Multi-task information bottleneck co-clustering for unsupervised cross-view human action categorization View paper
- [15] Incomplete multi-view data clustering with hidden data mining and fusion techniques View paper
- [16] Unifying complete and incomplete multi-view clustering through an information-theoretic generative model. View paper
- [17] Shared-private information bottleneck method for cross-modal clustering View paper
- [18] DMIB: Dual-Correlated Multivariate Information Bottleneck for Multiview Clustering. View paper
- [19] Multi-modal Spatial-modality Attentive Fusion for Studying Neuropsychiatric Disorders View paper
- [20] Information bottleneck fusion for deep multi-view clustering View paper
- [21] Differentiable Information Bottleneck for Deterministic Multi-View Clustering View paper
- [22] Shared-Neighbours methods for visual content structuring and mining View paper
- [23] GUARD: General Unsupervised Adversarial Robust Defense for Deep Multi-View Clustering via Information Bottleneck View paper
- [24] PHD THESIS Shared-Neighbours methods for visual content structuring and mining View paper
- [25] Dual global information guidance for deep contrastive multi-modal clustering View paper
- [26] Learning Compact Semantic Information for Incomplete Multi-View Missing Multi-Label Classification View paper
- [27] Calibrating Uncertainty Quantification of Multi-Modal LLMs using Grounding View paper
- [28] Pseudo-label calibration semi-supervised multi-modal entity alignment View paper
- [29] Calibrating multimodal learning View paper
- [30] Multimodal Emotion Recognition Calibration in Conversations View paper
- [31] Unveiling uncertainty: A deep dive into calibration and performance of multimodal large language models View paper
- [32] Does Adding a Modality Really Make Positive Impacts in Incomplete Multi-Modal Brain Tumor Segmentation? View paper
- [33] COLD fusion: Calibrated and ordinal latent distribution fusion for uncertainty-aware multimodal emotion recognition View paper
- [34] Calibrating class weights with multi-modal information for partial video domain adaptation View paper
- [35] DealMVC: Dual Contrastive Calibration for Multi-view Clustering View paper
- [36] Multi-Modal Learning with Bayesian-Oriented Gradient Calibration View paper
- [37] Debiased representation learning in recommendation via information bottleneck View paper
- [38] LM: Mutual Information Scaling Law for Long-Context Language Modeling View paper
- [39] Calibration bottleneck: Over-compressed representations are less calibratable View paper
- [40] Loss or gain: Hierarchical conditional information bottleneck approach for incomplete time series classification View paper
- [41] Learning Fair Graph Representations with Multi-view Information Bottleneck View paper
- [42] Analysis of Information Transfer Mechanism in Knowledge Distillation from an Information Theory Perspective View paper
- [43] Information theoretic counterfactual learning from missing-not-at-random feedback View paper
- [44] Estimating Information Flow in DNNs View paper
- [45] Scalable Mutual Information Estimation using Dependence Graphs View paper
- [46] DICE: Diversity in Deep Ensembles via Conditional Redundancy Adversarial Estimation View paper