Representation Learning with Conditional Information Flow Maximization

Dou $Hu^{1,2}$ and Lingwei Wei^{1*} and Wei Zhou¹ and Songlin $Hu^{1,2*}$

¹ Institute of Information Engineering, Chinese Academy of Sciences

² School of Cyber Security, University of Chinese Academy of Sciences

{hudou, weilingwei, zhouwei, husonglin}@iie.ac.cn

Abstract

This paper proposes an information-theoretic representation learning framework, named conditional information flow maximization, to extract noise-invariant sufficient representations for the input data and target task. It promotes the learned representations have good feature uniformity and sufficient predictive ability, which can enhance the generalization of pre-trained language models (PLMs) for the target task. Firstly, an information flow maximization principle is proposed to learn more sufficient representations for the input and target by simultaneously maximizing both inputrepresentation and representation-label mutual information. Unlike the information bottleneck, we handle the input-representation information in an opposite way to avoid the overcompression issue of latent representations. Besides, to mitigate the negative effect of potential redundant features from the input, we design a conditional information minimization principle to eliminate negative redundant features while preserve noise-invariant features. Experiments on 13 language understanding benchmarks demonstrate that our method effectively improves the performance of PLMs for classification and regression. Extensive experiments show that the learned representations are more sufficient, robust and transferable.

1 Introduction

The goal of deep representation learning (LeCun et al., 2015) is to transform the raw observational data into low-dimensional representations that are essential for various downstream tasks. In recent years, information-theoretic representation learning has been widely studied, aiming to discover useful representations in a principled manner. The InfoMax principle (Linsker, 1988) has extensive applications in the field of self-supervised representation learning (van den Oord et al., 2018; Hjelm

et al., 2019; Tschannen et al., 2020). In supervised scenarios, minimizing the standard cross-entropy is actually equivalent to maximizing the mutual information between the representations and the target task (Achille and Soatto, 2018; Boudiaf et al., 2020). But InfoMax tends to preserve potential redundant features that are irrelevant to the given target, leading to biased representations. Another noteworthy line of information-theoretic research is built upon the information bottleneck (IB) principle (Tishby et al., 1999; Tishby and Zaslavsky, 2015), which aims to discover compact and informative representations that can reduce redundant features from the inputs (Alemi et al., 2017; Belghazi et al., 2018; Fischer, 2020; An et al., 2023; Hu et al., 2024). IB seeks to find a maximally compressed representation of the input that preserves as much information as possible about the target, striking a balance between compression and prediction.

However, in the information flow (Goldfeld et al., 2019) of neural networks, directly reducing the mutual information between the input X and representations Z would violate the sufficiency constraint, and may lose the necessary information for the target task Y. Under the Markov chain constraint $Y \rightarrow X \rightarrow Z$, it's hard to determine beforehand how close we are to optimal compression, and this can easily lead to the over-compression issue of latent representations (Fischer, 2020; Hu et al., 2024). As a result, current IB-based methods would yield insufficient representations for the target task, and hamper prediction ability of neural networks.

To ensure sufficiency for the target task and mitigate the negative effect of redundant features, we propose a principled representation learning framework, named conditional information flow maximization (CIFM), to extract noise-invariant sufficient representations for the input and target. It promotes the learned representations have good feature uniformity and sufficient predictive ability, which can enhance the generalization of pre-trained

^{*}Corresponding author.

language models (PLMs) for the target task.

Firstly, we propose an information flow maximization (IFM) principle to learn more sufficient representations for the input and target. It simultaneously maximizes both input-representation and representation-label mutual information. Maximizing input-representation information I(X; Z) ensures sufficiency of the representations for the input X and preserves information relevant to the target Y, and maximizing representation-label information I(Y; Z) captures necessary information relevant to the target Y. In this way, the learned representation can be of good uniformity and sufficient predictive ability. Unlike IB that minimizes input-representation information, we handle the information in an opposite way to avoid the overcompression issue of latent representations.

Besides, we design a conditional information minimization (CIM) principle to mitigate the negative effect of potential redundant features from the input. In the information flow of $X \to Z$, InfoMax may introduces excessive and noisy information (Tian et al., 2020). For IFM, the task-irrelevant redundant nuisances features obtained by maximizing I(X; Z) interfere with the optimization of maximizing I(Y; Z). There are spurious correlations among these redundant features, forcing the model to learn a biased representation Z. As a conditional regularization term for IFM, the CIM principle eliminates negative redundant features while preserves noise-invariant features from inputs. Under the IFM principle with the conditional regularization, CIFM can extract noise-invariant sufficient representations for the input and target.

We conduct experiments on 13 natural language understanding benchmarks. The results demonstrate that CIFM can significantly improve the performance of PLMs for classification and regression. Our CIFM framework consistently achieves the best average performance compared to other methods, including 4 universal models and 7 representative deep representation learning technologies under different backbone models. For instance, with the RoBERTa backbone, CIFM improves average performance by +3.8% and +1.9% for classification and regression tasks compared to CE/MSE, respectively. Extended experiments prove that CIFM can enhance the model's generalization including out-of-distribution and data-constrained scenarios, robustness to random and adversarial noise, and transferability to new tasks. And the results also indicates that the learned representations by CIFM

are more sufficient, robust and transferable.

The contributions are as follows: 1) we propose an information flow maximization principle to learn more sufficient representations for the input and target by simultaneously maximizing both input-representation and representation-label information. 2) We design a conditional information minimization principle to eliminate negative redundant features while preserve noise-invariant features from the input. 3) We present an informationtheoretic CIFM framework to learn noise-invariant sufficient representations for the input and target. It can enhance the generalization of PLMs for better language understanding. 4) Experiments on 13 language understanding benchmarks demonstrate that CIFM achieves better performance under different backbone models. Extensive experiments show that the learned representations are more sufficient, robust and transferable.¹

2 Methodology

This section presents a new information-theoretic representation learning framework, named conditional information flow maximization (CIFM), to extract sufficient representations for the input data and target task, as well as eliminates negative redundant features from the input. It contains two principles including information flow maximization and conditional information minimization.

2.1 Information Flow Maximization

The IB principle can reduce redundant features from the inputs by finding a compressed representation of the input that maximally preserves information about the output. However, in the information flow (Goldfeld et al., 2019) of neural networks, directly minimizing the mutual information between the input X and representations Z would violate the sufficiency constraint, and may lose the necessary information for the target task Y. Fischer (2020) and Hu et al. (2024) have shown that the compression term I(X; Z) would reduce necessary information related to target task Y under the Markov chain constraint $Y \rightarrow X \rightarrow Z$. As a result, current IB-based methods would yield insufficient representations for the target task, and hamper prediction ability of neural networks.

As shown in Figure 1, regions being minimized by I(X; Z) overlap with relevancy regions being

 $^{^{\}rm l} The source code is available at https://github.com/zerohd4869/CIFM$



Figure 1: Venn information diagram comparison of our CIFM with existing principles. The learned representations by each principle is circled by the red dashed line.

I

maximized by I(Z; Y). To preserve informative signals from the input X relevant to the target Y (i.e., region (1)) as much as possible in the information flow of $X \to Z$, we use the idea of Info-Max to maximize I(X; Z). Besides, we also maximize I(Y; Z) to capture necessary information relevant to the target task Y. Combing the above two terms, the principle of information flow maximization (IFM) simultaneously maximizes both input-representation and representation-label mutual information. It can be formulated as the maximization of the following Lagrangian,

$$\max I(Y;Z) + \beta I(X;Z), \tag{1}$$

subject to the Markov chain constraint, i.e., $Y \rightarrow X \rightarrow Z$. β is a parameter that balances the tradeoff between the informativeness for X and Y.

In Equation (1), the second term that maximizing input-representation information I(Y; Z) ensures sufficiency of the representations for the input X and preserves informative signals from X relevant to the target Y. The first term that maximizing representation-label information I(Y; Z)captures necessary information relevant to the target Y. In this way, the learned representation by IFM can be of good uniformity and sufficient predictive ability. Unlike the IB principle that minimizes input-representation information, we handle the information in an opposite way to avoid the over-compression issue of latent representations.

Implementation of IFM The implementation of IFM contains two terms, i.e., maximizing I(Y; Z) and I(X; Z). First, we maximize the lower bound of I(Y; Z) by estimating the conditional entropy of the target Y given representations Z. Following Kolchinsky et al. (2019) and Hu et al. (2024), we use cross-entropy (CE) and mean squared error (MSE) as the estimators for classification and regression, respectively, i.e., $-\log \text{Softmax}(z_i, y_i)$ and $||z_i - y_i||^2$ for the *i*-th sample. Then, maximizing I(X; Z) can be optimized by using the mutual

information estimators, e.g., $InfoNCE^2$ (van den Oord et al., 2018) and MINE (mutual information neural estimator) (Belghazi et al., 2018).

Here, we take InfoNCE as the default estimator for maximizing I(X; Z). According to the information flow of $Z \leftarrow X \rightarrow Z'$, the Markov chain rule states that $I(X; Z) \ge I(Z; Z')$. Maximizing the lower bound of I(X; Z), i.e., I(Z; Z'), can be optimized by using InfoNCE estimator. Specifically, for a sample x_i with the representation z_i , its positive key is the augmented sample obtained by dropout, and its negative keys are the other samples in the batch N. In this case, we have,

$$(X;Z) \ge I(Z;Z') \approx \log(K) - \frac{1}{N} \sum_{i=1}^{N} \log \frac{\exp(z_i \cdot z_i^+ / \tau)}{\sum_{j=1}^{K} \exp(z_i \cdot z_j^- / \tau)},$$
 (2)

where z_i^+ is the positive key, which is the augmented sample obtained by dropout. z_j^- represents the negative keys, which are the representations of the other samples in the batch. K is the number of negative keys, i.e., 2N - 2. $\tau > 0$ is a scalar temperature parameter that controls the concentration or separation of probability distribution.

Alternatively, we also present how to utilize the MINE estimator to maximize I(X; Z). Specifically, we use a neural network to approximate the lower bound of I(X; Z), i.e.,

$$I(X;Z) \ge I_{\Omega}(X;Z)$$

= sup $\mathbb{E}_{\mathbb{P}_{XZ}}[T_{\omega}(Z|X)] - \log(\mathbb{E}_{\mathbb{P}_{X}\otimes\mathbb{P}_{Z}}[e^{T_{\omega}(Z|\tilde{X})}]),$ (3)

where \mathbb{P}_{XZ} is the joint distribution and $\mathbb{P}_X \otimes \mathbb{P}_Z$ is the product of the marginals. T_{ω} is a function parametrized by an MLP (i.e., a fully-connected neural network with two hidden layers) with the parameter $\omega \in \Omega$. The expectations of $I_{\Omega}(X; Z)$ are estimated by shuffling the samples from the joint distribution along the batch axis. Additionally,

²InfoNCE is a softmax-based version of noise-contrastive estimation (Gutmann and Hyvärinen, 2010).

an exponential moving average is used to make a biased estimate.

2.2 Conditional Information Minimization

In the information flow of $X \rightarrow Z$, InfoMax principle may introduce excessive and noisy information (Tian et al., 2020). In our IFM, the task-irrelevant redundant nuisances features obtained by maximizing I(X; Z) interfere with the optimization of maximizing I(Y; Z). There are spurious correlations among these redundant features, forcing the model to learn a biased representation Z.

To address the issue, we design a new conditional information minimization (CIM) term to eliminate negative redundant features while preserve noise-invariant features from inputs. As shown in Figure 1, we categorize all redundant information (i.e., region (2) and (3)) into two broad categories: negative (i.e., region (3)) and nonnegative (i.e., region (2)), depending on whether it impedes the prediction ability of the target task. The negative features mean the nuisances that exist spurious correlations among redundant information and have a negative impact on task prediction. These non-negative features may encompass some inherent structured characteristics in the data, and preserving them during the learning process contributes to improving the model's generalization.

Given the target Y, we define the conditional effective mutual information between the input X and the representations Z as the subset of effective information that is irrelevant to the target Y. Here, effective information (Achille et al., 2019; Terzi et al., 2021) in the activations Z_{δ} can be seen as the amount of information about the input X that is retained after adding the noise δ . The CIM principle can be formulated as the minimization of the conditional effective information between the input X and the activations Z_{δ} given the target Y, i.e.,

$$\min I(X; Z_{\delta}|Y), \tag{4}$$

where $Z_{\delta} = \phi_{w+\delta}(X)$ are the activations of the intermediate layer computed by the perturbed weights $w + \delta$. Under the target Y, the higher the noise level $\delta (|| \delta ||_2 \le \epsilon)$, the effective information in the activations Z_{δ} about X gradually decreases to the optimum (i.e., region (2) + (3) in Figure 1) with the maximum noise level ϵ .

Implementation of CIM Based on the relation between Fisher information and effective information (Achille et al., 2019), we use the Fisher Information of the weights to control the lower bound of the conditional effective information:

$$I(X; Z_{\delta}|Y) \approx H(X|Y) - \mathbb{E}_{x,y} \left[\frac{1}{2} \log \left(\frac{(2\pi e)^k}{|F_{z|x,y}|} \right) \right],$$
(5)

where $F_{z|x,y}$ represents the Fisher information in the activations Z_{δ} given the input X and target Y, i.e., $F_{z|x,y} = \mathbb{E}_z[\nabla_x^2 \log p(z|x,y)]$. H(X|Y) is the conditional entropy of input X given target Y. To reduce the conditional effective information between inputs X and activations Z_{δ} given the target Y, it is sufficient to decrease the Fisher $|F_{z|x,y}|$, that is, increasing ϵ while keeping the target Y of the noise features remains the same.

Inspired by Terzi et al. (2021), we use gradientbased adversarial training (Goodfellow et al., 2015; Miyato et al., 2017) to approximately estimate the minimization of the Fisher information $F_{z|x,y}$ in the activations Z_{δ} given the input X and target Y. Formally, denote (x, y) as a mini-batch input sampled from distribution D and $p(y|x;\theta)$ as a model with the parameter θ . The worst-case perturbation δ under the given Y can be computed by the back-propagation in the network, i.e., $\max_{\|\delta\|_q \le \epsilon} KL(p(\cdot|x;\theta_{\delta});p(\cdot|x;\theta))$ with a Lagrange constraint $\mathcal{L}_{IFM}(x, y; \theta_{\delta})$ and an L_q norm constraint. At each step of training, we identify the adversarial perturbations δ against the current network with the parameter θ , and put them on the weights of first several hidden layers of the network. When using a Transformer-based network (e.g., RoBERTa) as the model backbone, we probabilistically put the perturbations on the weights of an embedding layer and the first hidden layer of the Transformer based on a predefined perturbation rate. With a linear approximation (Goodfellow et al., 2015), an L_q norm-ball constraint, a certain radius ϵ for δ , and the original objective \mathcal{L}_{IFM} , the formulation of adversarial objective \mathcal{L}_{CIM} is,

$$\min_{\theta} \mathbb{E}_{(x,y)\sim D} \max_{\|\delta\|_2 \le \epsilon} \mathcal{L}_{IFM}(x,y;\theta_{\delta}),$$
(6)
where $\delta = -\epsilon g / \|g\|_q, \ g = \nabla_x \log p(y|x;\hat{\theta}).$

Under the objective of \mathcal{L}_{IFM} , the conditional information minimization loss \mathcal{L}_{CIM} can be viewed as an adversarial regularizer.

2.3 CIFM Framework

We incorporate the conditional information minimization into information flow maximization, named conditional information flow maximization (CIFM). The total optimization principle can be,

$$\max I(Y;Z) + \beta I(X;Z) - I(X;Z_{\delta}|Y).$$
(7)

The first two terms simultaneously maximize both input-representation and representation-label mutual information, which can learn more sufficient representations for the input and target. The last term eliminates negative redundant features while preserves noise-invariant features from the input to mitigate the negative effect of potential redundant features. Totally, CIFM can extract noise-invariant sufficient representations for the input and target. It promotes the learned representations have good uniformity and sufficient predictive ability, which can enhance the generalization of pre-trained language models for the target task.

As shown in Figure 1, given the target task, Info-Max maximizes information between the representations Z and the target task Y (e.g., the CE loss for classification), which preserves a lot of potential redundant features for the target task. There are spurious correlations among these redundant features (i.e., region (3)), forcing the model to learn biased representations Z. IB is prone to learning over-compressed representations due to the challenge in balancing between compression and task prediction. Directly minimizing the mutual information between the input X and representations Z can easily reduce necessary information (i.e., region (I)) for the target task Y to some extent under the Markov chain constraint. Different from them, CIFM finds more sufficient representations (i.e., region (1) and (2)) for the input and target, as well as adversarially eliminates negative redundant features (i.e., region (3)) from the input.

3 Experiments

3.1 Experimental Setups

Downstream Tasks and Datasets We conduct experiments on a variety of classification and regression tasks. Specifically, we utilize 10 classification benchmarks: EmojiEval (Barbieri et al., 2018) for emoji prediction, HatEval (Basile et al., 2019) for hate speech detection, IronyEval (Hee et al., 2018) for irony detection, OffensEval (Zampieri et al., 2019) for offensive language detection, SentiEval (Rosenthal et al., 2017) for sentiment analysis, StanceEval (Mohammad et al., 2016) for stance detection, and 4 emotion-related benchmarks from different domains (i.e., EmotionEval (Mohammad et al., 2018), ISEAR (Scherer and Wallbott, 1994), MELD (Poria et al., 2019), and GoEmotions (Demszky et al., 2020)) for categorical emotion analysis. For regression, we use 3 benchmarks: STS-B

(Cer et al., 2017) for semantic similarity prediction, *CLAIRE* (Roth et al., 2022) for plausible clarification ranking, and *EmoBank* (Buechel and Hahn, 2017) for dimensional emotion analysis. More descriptions can be found in Appendix B.1.

Comparison Methods We compare against the 4 universal models (i.e., SVM (Cortes and Vapnik, 1995), FastText (Joulin et al., 2017), BiLSTM (Hochreiter and Schmidhuber, 1997), and GPT- (3.5^3) and 7 representative deep representation learning technologies (i.e., CE/MSE, CE+CP (Pereyra et al., 2017), CE/MSE+AT (Miyato et al., 2017), CE+SCL (Gunel et al., 2021), VIB (Alemi et al., 2017; Mahabadi et al., 2021), MINE-IB (Belghazi et al., 2018) and MEIB (An et al., 2023)) with 2 different backbone models. For these representation learning technologies, we use two PLMs, i.e., BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), as the backbone models for finetuning on downstream tasks. Concretely, we use *bert-base-uncased*⁴ and *roberta-base*⁴ to initialize BERT and RoBERTa for fine-tuning on downstream tasks. See Appendix B.2 for more details.

Evaluation Metrics We use the same evaluation metric from the original tasks. For classification tasks, the macro-averaged F1 over all classes is applied in most cases. There are three exceptions: stance (macro-averaged of F1 of favor and against classes), irony (F1 of ironic class), and sentiment analysis (macro-averaged recall). For regression tasks, we use both Pearson and Spearman correlation coefficients on STS-B and CLAIRE, and Pearson correlation for each VAD dimension on EmoBank. We also report a global metric based on the average of all dataset-specific metrics, i.e., $\frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_t}\sum_{n=1}^{N_t}M_{t,n}$. Here, $M_{t,n}$ denotes the performance for the *n*-th metric in the *t*-th task, N_t denotes the number of metrics in the *t*-th task, and T refers to the number of tasks. Besides, for the target dataset, the t-test (Kim, 2015) is used to verify the statistical significance of the differences between the results of our CIFM and the best non-CIFM methods using the same backbone.

Implementation Details All experiments are conducted on a single NVIDIA Tesla A100 80GB card. The validation sets are used to tune hyperparameters and choose the optimal model. For each method, we run five random seeds and re-

³https://chat.openai.com

⁴https://huggingface.co/

Mathada	EmojiEval	EmotionEval	HatEval	IronyEval	OffensEval	SentiEval	StanceEval	ISEAR	MELD	GoEmotions	Ave
Methous	M-F1	M-F1	M-F1	F1(i.)	M-F1	M-Recall	M-F1 (a. & f.)	M-F1	M-F1	M-F1	Avg.
SVM [†]	29.30	64.70	36.70	61.70	52.30	62.90	67.30	-	-	-	-
FastText [†]	25.80	65.20	50.60	63.10	73.40	62.90	65.40	-	-	-	-
$BiLSTM^{\dagger}$	24.70	66.00	52.60	62.80	71.70	58.30	59.40	-	-	-	-
GPT-3.5	6.34 ± 0.01	$73.23{\scriptstyle \pm 0.18}$	$48.30{\scriptstyle \pm 0.11}$	66.81±3.26	63.71 ± 0.13	$40.40{\scriptstyle \pm 3.13}$	39.45 ± 0.10	$67.22{\scriptstyle\pm0.09}$	$41.46{\scriptstyle\pm0.11}$	$25.21{\scriptstyle\pm0.08}$	47.21
BERT baci	kbone										
CE	22.30 ± 0.60	76.05 ± 1.41	44.67 ± 1.78	59.38 ± 3.01	80.16±1.26	70.54 ± 0.44	65.21 ± 0.71	67.17 ± 0.78	39.80 ± 0.84	46.29 ± 0.79	57.16
CE+CP	21.91 ± 0.71	76.28 ± 1.20	$45.97{\scriptstyle\pm2.93}$	$64.06{\scriptstyle\pm2.41}$	78.99 ± 1.57	$70.68{\scriptstyle\pm0.31}$	65.83±0.39	67.20 ± 0.95	39.54 ± 1.61	46.39 ± 0.63	57.69
CE+AT	22.93 ± 0.70	75.08 ± 1.23	$46.30{\scriptstyle \pm 3.61}$	$64.23{\scriptstyle\pm2.04}$	79.68 ± 1.59	$70.55{\scriptstyle \pm 0.57}$	66.46±1.13	65.70 ± 0.69	$39.84{\scriptstyle\pm0.38}$	$47.37{\scriptstyle\pm0.54}$	57.81
CE+SCL	21.72 ± 0.51	75.43 ± 1.37	45.86 ± 1.15	$65.39{\scriptstyle \pm 2.46}$	80.20 ± 0.56	70.70 ± 0.79	65.34 ± 0.60	67.54 ± 0.64	40.00 ± 1.96	46.50 ± 0.46	57.87
VIB	21.31 ± 0.62	77.37 ± 0.71	45.99±1.93	63.82 ± 1.00	80.37 ± 1.11	$70.39{\scriptstyle \pm 0.31}$	65.43 ± 0.60	67.24 ± 0.57	38.52 ± 0.51	45.89 ± 1.10	57.63
MINE-IB	21.29 ± 0.31	76.60 ± 0.41	47.64 ± 2.11	$65.86{\scriptstyle \pm 2.57}$	$78.67{\scriptstyle\pm2.28}$	$69.85{\scriptstyle \pm 0.54}$	65.35 ± 0.88	67.62 ± 0.40	41.23 ± 0.67	46.87 ± 0.42	58.10
MEIB	21.87 ± 0.73	76.70 ± 0.82	48.27 ± 1.72	$65.87{\scriptstyle\pm2.14}$	$80.49{\scriptstyle \pm 0.81}$	$70.55{\scriptstyle \pm 0.57}$	65.59 ± 1.58	$67.44{\scriptstyle\pm0.50}$	$39.30{\scriptstyle \pm 0.61}$	$46.26{\scriptstyle \pm 0.81}$	58.23
CIFM	$24.28^*{\scriptstyle\pm0.39}$	77.74 ± 1.07	$59.22^*{\scriptstyle \pm 1.75}$	65.51 ± 1.61	$80.67{\scriptstyle\pm0.99}$	$70.66{\scriptstyle \pm 0.28}$	$67.81^* \pm 0.97$	$71.51^*{\scriptstyle\pm0.51}$	$41.72^*{\scriptstyle\pm0.80}$	$48.42^*{\scriptstyle\pm0.74}$	60.75
RoBERTa	backbone										
CE	30.25 ± 1.32	77.41±1.33	45.49 ± 4.70	57.99±4.96	78.74 ± 2.20	71.80±0.93	66.78±1.34	70.00 ± 0.45	39.23 ± 0.41	46.64 ± 1.15	58.43
CE+CP	31.12 ± 0.84	77.54 ± 0.70	$48.59{\scriptstyle\pm3.28}$	58.75±6.19	79.50 ± 0.98	72.82 ± 0.29	66.89 ± 1.67	70.58 ± 0.71	40.74 ± 0.89	47.98 ± 0.65	59.45
CE+AT	32.00 ± 0.93	77.30 ± 1.07	44.71 ± 4.76	60.17 ± 3.17	79.81 ± 1.11	$72.51{\scriptstyle \pm 0.44}$	67.81 ± 0.95	$70.97{\scriptstyle\pm0.68}$	$40.10{\scriptstyle \pm 0.60}$	$47.89{\scriptstyle\pm1.21}$	59.33
CE+SCL	31.09 ± 1.85	76.98 ± 2.02	49.51 ± 2.86	60.71±4.23	$80.39{\scriptstyle \pm 0.83}$	73.16 ± 0.44	66.73±1.54	70.26 ± 0.45	40.64 ± 1.02	$47.87{\scriptstyle\pm0.86}$	59.72
VIB	29.71 ± 0.79	77.99 ± 0.86	$49.39{\scriptstyle\pm3.08}$	$59.93{\scriptstyle\pm4.57}$	$79.63{\scriptstyle \pm 0.66}$	72.81 ± 0.39	68.40 ± 0.52	70.74 ± 0.44	38.94 ± 0.55	46.23 ± 0.18	59.38
MINE-IB	31.70 ± 0.45	$78.79{\scriptstyle \pm 0.58}$	$46.39{\scriptstyle\pm2.82}$	$57.39{\scriptstyle\pm8.27}$	$79.76{\scriptstyle \pm 0.67}$	$72.85{\scriptstyle \pm 0.56}$	67.27 ± 1.00	$70.15{\scriptstyle \pm 0.58}$	$41.80{\scriptstyle \pm 2.14}$	$48.88{\scriptstyle\pm1.04}$	59.50
MEIB	29.94 ± 1.30	$78.73{\scriptstyle \pm 0.90}$	$49.34{\scriptstyle\pm2.42}$	$60.54{\scriptstyle\pm2.70}$	$79.68{\scriptstyle \pm 0.98}$	$72.78{\scriptstyle \pm 0.29}$	67.89 ± 1.70	$70.86{\scriptstyle \pm 0.61}$	$39.00{\scriptstyle\pm0.37}$	47.18 ± 1.15	59.59
CIFM	$32.32{\scriptstyle\pm0.87}$	$\textbf{79.63}^{*} \pm 0.57$	$59.55^*{\scriptstyle\pm0.91}$	63.02*±3.81	$80.96^{*} \pm 0.55$	$72.93{\scriptstyle\pm0.27}$	$69.01^* \pm 0.80$	$\textbf{71.66}^{*} \pm 0.68$	$43.99^{*} \pm 1.10$	$\textbf{49.51}^{*}{\scriptstyle\pm0.31}$	62.26

Table 1: Classification evaluation (%) on 10 benchmarks. BERT and RoBERTa are the backbone models for fine-tuning on each benchmark. [†] means the results are from Barbieri et al. (2020). For other methods, we run five random seeds and report the average result on test sets. Best results for each benchmark are highlighted in bold. * represents statistical significance over state-of-the-art scores for the same backbone under the *t*-test (p < 0.05).

Methods	ST	S-B	CLA	AIRE		EmoBank		A.v.a
Methous	Spearman	Pearson	Spearman Pearson		Pearson (v)	Pearson (a)	Pearson (d)	Avg.
MSE	88.33±0.32	$88.80{\scriptstyle\pm0.36}$	50.37±5.90	49.10±5.74	$80.62{\scriptstyle \pm 0.64}$	55.43±1.95	$49.51_{\pm 1.64}$	66.72
MSE+AT	88.40 ± 0.50	$89.01{\scriptstyle \pm 0.37}$	$53.09{\scriptstyle\pm0.64}$	$51.87{\scriptstyle\pm0.65}$	$81.04{\scriptstyle \pm 0.82}$	56.73±1.16	$51.48{\scriptstyle \pm 1.46}$	68.15
VIB	88.45 ± 0.50	$89.01{\scriptstyle \pm 0.40}$	$52.86{\scriptstyle\pm0.88}$	$51.66{\scriptstyle \pm 0.78}$	$79.41{\scriptstyle \pm 0.82}$	$55.07{\scriptstyle\pm0.96}$	46.50 ± 3.85	67.10
MEIB	88.61±0.14	$89.13{\scriptstyle \pm 0.17}$	$52.85{\scriptstyle\pm0.72}$	$51.39{\scriptstyle\pm0.81}$	$79.32{\scriptstyle \pm 0.63}$	55.58 ± 1.15	$46.63_{\pm 3.10}$	67.17
CIFM	$88.94^{*}{\scriptstyle \pm 0.38}$	$89.52^{\ast}{\scriptstyle \pm 0.32}$	$53.62^{\ast}{\scriptstyle \pm 0.51}$	$52.88^{\ast}{\scriptstyle \pm 0.30}$	$81.64^{\ast}{\scriptstyle\pm0.84}$	$57.06^{\ast}{\scriptstyle \pm 0.95}$	$51.60{\scriptstyle\pm1.11}$	68.64

Table 2: Regression evaluation (%) on 3 benchmarks with RoBERTa backbone. * represents statistical significance over state-of-the-art scores under the *t*-test (p < 0.05).

Mathada	Avg. Classification	Avg. Regression
Methous	BERT / RoBERTa	RoBERTa
CIFM	60.75 / 62.26	68.64
- w/o CIM	58.96 / 60.96	66.77
- w/o CIM & IFM	57.27 / 57.92	67.32

Table 3: Ablation results (%). We report the average results on 10 classification benchmarks and 3 regression benchmarks. See Appendix C.1 for the detailed results.

port the average result of the test sets. The network parameters are optimized by using Adamax optimizer (Kingma and Ba, 2015) with the learning rate of $5e^{-5}$, the weight decay coefficient of $\{0, 0.01, 0.001\}$. More implementation details are listed in Appendix B.3.

3.2 Overall Results

The overall results for 10 classification benchmarks and 3 regression benchmarks are summarized in Table 1 and Table 2, respectively. Our CIFM consistently obtains the best average performance over comparison methods. When using RoBERTa, CIFM improves average performance by +3.8% and +1.9% for classification and regression tasks compared to CE/MSE, respectively. The results indicate the good generalization ability of our method to unseen test sets and show the superiority of our method. We notice that CIFM achieves big improvements for *HatEval* and *IronyEval*, i.e., +14.1% macro-F1 scores and +5.0% F1 scores of ironic. In *HateEval*, there is a phenomenon of disparity in topic distribution between the validation set and the test set. Besides, the task of *IronyEval* requires complex semantic understanding, and the differences between ironic and non-ironic texts are usually subtle. The results indicate that CIFM has a good generalization capability on the above scenarios, i.e., topic shifts and subtle semantic labels.

3.3 Ablation Study

Loss Analysis Table 3 shows an ablation study by removing the conditional information minimization (w/o CIM) and information flow maximization (w/o IFM). From results, the full CIFM achieves the best performance on classification and regression tasks. It proves the effectiveness of combing



Figure 2: Comparison results of CIFM with different MI Estimators and the CE baseline on classification tasks. RoBERTa is the default backbone model.

Mathada	EmotionEval	ISEAR	MELD	4.00	
Methods	\rightarrow GoEmotions	\rightarrow GoEmotions	\rightarrow GoEmotions	Avg.	
CE	73.79±2.57	42.99±2.10	30.71±0.54	49.16	
CE+AT	72.54±3.89	44.11±1.44	32.05±1.69	49.57	
VIB	74.73±3.52	41.88±1.65	30.50±1.05	49.03	
MEIB	75.55±2.05	42.10±0.61	30.11±1.33	49.25	
CIFM	75.88±0.93	58.85±1.93	36.44±1.76	57.06	

Table 4: Out-of-distribution evaluation results (%). For example, $MELD \rightarrow GoEmotions$ refers to training the model on training set of MELD and predicting with the test set of GoEmotions. We experiment with RoBERTa backbone. We run five random seeds and report the average results on test sets of target domains. Labels that do not appear in the training corpus are not evaluated.

CIM and IFM. Specifically, learning more sufficient representations by IFM can boost the performance of classification tasks. Mitigating the negative effect of redundant features by CIM is more helpful for regression.

Effectiveness Evaluation with Different MI Estimators We further implement our CIFM by replacing the default InfoNCE estimator with MINE (Belghazi et al., 2018) to estimate I(X, Z). Figure 2 shows comparison results of CIFM with different mutual information (MI) estimators and the CE baseline on the classification benchmarks. From results, our CIFM and its variant CIFM w/ MINE consistently achieve better results on all tasks. It confirms the effectiveness of our CIFM principle under different MI estimators.

3.4 Generalization Evaluation

We evaluate the generalization of CIFM under the following two settings, i.e., out-of-distribution (OOD) and data-constrained scenarios.

Out-of-Distribution Generalization Evaluation We choose emotion-related benchmarks (e.g., *EmotionEval*, *ISEAR*, *MELD*, and *GoEmotions*), which are collected from different domains and aim to predict the label of different emotion taxonomies.

To implement OOD scenarios, we train the model on the original training set from a source domain, select the best model based on the validation set of the source domain, and test on the test set of a target domain. To avoid the interference of label mapping bias between different taxonomies, each model is trained on the dataset with coarse-grained taxonomy to predict the label for another dataset with fine-grained taxonomy (i.e., GoEmotions). Table 4 shows the performance under OOD scenarios. CIFM obtains the best results on all OOD settings. Comparing to CE, CIFM achieves +7.9% improvements in terms of average scores. This results exhibit that CIFM can promote the learning of more sufficient representations for target task than others, which enhances the generalization in handling OOD scenarios across different domain shifts.

Evaluation under Data-constrained Scenarios

We experiment under different ratios of the training set to evaluate the generalization when training with data-constrained scenarios. Specifically, given a predefined ratio (e.g., 20%) and a random seed, we randomly sample from the original training set. We obtain 5 training subsets by independently and repeatedly sampling five times from the original training set with 5 different random seeds. All methods are trained on these 5 subsets of the training set, and we report the average results on the test set. Figure 3 shows results of CE, CE+AT, VIB, MEIB, and our CIFM against different sizes of training set with RoBERTa backbone. With a smaller ratio, the comparison methods struggle to capture complex patterns from the limited data, resulting in poor generalization on the test sets. Our CIFM achieves superior performance on all classification tasks against different ratios of the training set. This proves that CIFM can enhance the model's generalization ability, even under conditions of limited training data.

3.5 Robustness Evaluation

We experiment to evaluate the model's robustness to noise samples under various optimization objectives during training. We adjust different strengths of random and adversarial perturbations on the test set. The random perturbations are from a multivariate Gaussian, and the adversarial perturbations are produced by a fast gradient method (Miyato et al., 2017). These perturbations are scaled by the L_2 norm and then applied to the embedding layer in the testing process. Following the empirical robust



Figure 3: Results of different methods against different sizes of training set with RoBERTa backbone.



Figure 4: Robust scores against different random perturbation strengths. RoBERTa is the default backbone.



Figure 5: Robust scores against different adversarial perturbation strengths. RoBERTa is the default backbone.

evaluation (Carlini and Wagner, 2017; Hu et al., 2023), we report the robust scores in terms of original evaluation metrics on noise samples generated from original test sets for each benchmark.

From Figure 4 and Figure 5, CIFM gains better robust scores over other objectives on all classification tasks. Compared to CE, CIFM achieves an average increase of +3.2% and +3.6% in robust scores under random and adversarial noise, respectively. It indicates that CIFM can safely discard the specific redundancy regions (i.e., region ③ in Figure 1) that are negative to the target task. That is, CIFM learns noise-invariant features from input, which could be beneficial to noise scenes.

3.6 Transferability Evaluation

Unlike IB-based methods that learn representations by eliminating all target-irrelevant information, CIFM adversarially preserves noise-invariant but target-irrelevant features from the input, i.e., region (2) in Figure 1. We believe that although the information of this region may not directly be relevant to the target, it can be valuable in learning more sufficient representations for the input data, accordingly to enhance the model's transferability. To prove this, we evaluate both linear and nonlinear transferability to new tasks. Two large-scale GoEmotions and SentiEval datasets are selected as the source dataset for pre-training. After pre-training, we freeze parameters of feature extractor (including parameters of an embedding layer and all hidden layers of the Transformer) and fine-tune a new classifier for each target task. We implement linear and nonlinear transferability evaluations by employing a linear hidden layer (Linear) and a convolutional neural network (CNN) as classifiers, respectively.

l	Methods			SentiEval \rightarrow Others								
Cls.	Pre-train Obj.	EmojiEval	EmotionEval	HatEval	IronyEval	OffensEval	SentiEval	StanceEval	ISEAR	MELD	Avg.	Avg.
	CE	5.89	59.39	58.52	49.36	74.86	66.42	38.47	49.58	31.57	48.23	37.88
	CE+AT	5.74	59.65	58.47	46.23	74.67	65.86	37.43	48.63	31.62	47.59	38.20
Linear	VIB	5.64	52.51	58.01	45.54	72.57	64.87	39.05	47.11	30.00	46.14	37.07
	MEIB	4.44	51.10	57.64	44.67	71.38	62.79	37.89	47.98	28.61	45.17	35.26
	CIFM	8.10	59.74	59.33	47.96	75.41	67.03	38.41	50.17	31.71	48.65	40.44
-	CE	11.43	72.33	45.40	59.35	75.84	70.22	56.89	61.88	41.29	54.96	55.13
	CE+AT	10.74	73.03	49.16	57.36	78.30	70.49	55.29	61.93	41.20	55.28	55.40
CNN	VIB	12.52	72.02	48.52	57.83	75.43	69.62	55.91	60.38	40.06	54.70	53.07
	MEIB	10.03	71.45	47.09	51.39	75.10	68.43	47.33	58.93	38.55	52.03	50.94
	CIFM	15.19	75.30	50.58	57.14	78.64	70.99	58.80	64.71	41.37	56.97	56.21

Table 5: Transferability results (%) of *GoEmotions* \rightarrow *Others* and *SentiEval* \rightarrow *Others* with RoBERTa backbone. *GoEmotions* and *SentiEval* are the source datasets for pre-training. Pre-train Obj. indicates the pre-training objective. Cls. refers to the classifier when fine-tuning. Best results for linear (i.e., Linear) and nonlinear (i.e., CNN) classifiers are highlighted in bold. Detailed results of *SentiEval* \rightarrow *Others* are listed in Appendix C.2.



Figure 6: Quality analysis of the learned representations by different optimization objectives. The X-axis refer to the uniformity of hidden representations, and the Y-axis refer to the ARI score of output representations. The lower uniformity means the better sufficiency for the input, and the higher ARI means the better sufficiency for the target.

Table 5 summarizes the average scores over the other nine target tasks against two types of classifiers. Notably, the pre-trained CIFM consistently achieves the best average scores on new tasks by fine-tuning linear and non-linear classifiers. The results demonstrate that CIFM can capture more sufficient representations from the input and the learned representations are more transferable. Besides, employing a CNN classifier tends to yield higher scores compared to a linear classifier. This proves nonlinear classifiers are instrumental in enhancing the model's adaptability across tasks.

3.7 Representation Quality Evaluation

To evaluate the quality of the learned representations, we use two metrics to measure the sufficiency of both the input and target task. Following Wang and Isola (2020), we choose the uniformity metric to measure the preserved maximal information of hidden representations from the input. Following Hu et al. (2023), we apply the adjusted rand index (ARI) metric to assess the preserved maximal information of output representations for label structure. The lower the uniformity, the better sufficiency for the data. Conversely, a higher ARI indicates a better sufficiency for the target.

Figure 6 shows uniformity and ARI of representations learned by various optimization objectives. CIFM not only extracts more uniform distributions on the hidden feature space, but also achieves better supervised clustering ability on the output representations. This implies that our CIFM exhibits more sufficient representations for the input and target, and the learned representations have good feature uniformity and sufficient predictive ability.

4 Conclusion

We propose an information-theoretic CIFM framework to extract noise-invariant sufficient representations for the input data and target task. Firstly, an IFM principle is proposed to learn more sufficient representations for the input and target. Besides, a CIM principle is designed to eliminate negative redundant features while preserve noise-invariant features from the input. Experiments on 13 language understanding benchmarks demonstrate that CIFM effectively improves the performance of PLMs for classification and regression. Extensive experiments show that the learned representations are more sufficient, robust and transferable.

Limitations

This paper introduces a new information-theoretic framework to extract sufficient representations for the input data and target task, as well as eliminates negative redundant features from the input. It promotes learned representations with good uniformity and sufficient predictive ability, thereby enhancing the generalization of pre-trained language models for the target task. However, the technique is studied only for different classification and regression tasks; its effectiveness in generation tasks remains to be further validated. Additionally, the principles proposed in this paper require more in-depth theoretical investigations.

Acknowledgements

This work was supported by the National Key Research and Development Program of China (No. 2022YFC3302102), and the Postdoctoral Fellowship Program of CPSF (No. GZC20232969). The authors thank the anonymous reviewers and the meta-reviewer for their helpful comments.

References

- Alessandro Achille, Giovanni Paolini, and Stefano Soatto. 2019. Where is the information in a deep neural network? *arXiv preprint arXiv:1905.12213*.
- Alessandro Achille and Stefano Soatto. 2018. Information dropout: Learning optimal representations through noisy computation. *IEEE TAPAMI*, 40(12):2897–2905.
- Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. 2017. Deep variational information bottleneck. In *ICLR (Poster)*.
- Sungtae An, Nataraj Jammalamadaka, and Eunji Chong. 2023. Maximum entropy information bottleneck for uncertainty-aware stochastic embedding. In *CVPR Workshops*, pages 3809–3818.
- Francesco Barbieri, José Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. In *EMNLP (Findings)*, pages 1644–1650.
- Francesco Barbieri, José Camacho-Collados, Francesco Ronzano, Luis Espinosa Anke, Miguel Ballesteros, Valerio Basile, Viviana Patti, and Horacio Saggion. 2018. Semeval 2018 task 2: Multilingual emoji prediction. In SemEval@NAACL-HLT, pages 24–33.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019.

Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In *SemEval@NAACL-HLT*, pages 54–63.

- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, R. Devon Hjelm, and Aaron C. Courville. 2018. Mutual information neural estimation. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 530–539.
- Malik Boudiaf, Jérôme Rony, Imtiaz Masud Ziko, Eric Granger, Marco Pedersoli, Pablo Piantanida, and Ismail Ben Ayed. 2020. A unifying mutual information view of metric learning: Cross-entropy vs. pairwise losses. In *ECCV* (6), volume 12351 of *Lecture Notes in Computer Science*, pages 548–564. Springer.
- Sven Buechel and Udo Hahn. 2017. Emobank: Studying the impact of annotation perspective and representation format on dimensional emotion analysis. In *EACL* (2), pages 578–585. Association for Computational Linguistics.
- Nicholas Carlini and David A. Wagner. 2017. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, pages 39–57. IEEE Computer Society.
- Daniel M. Cer, Mona T. Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In SemEval@ACL, pages 1–14.
- Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, pages 2172–2180.
- Corinna Cortes and Vladimir Vapnik. 1995. Supportvector networks. *Machine learning*, 20:273–297.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan S. Cowen, Gaurav Nemade, and Sujith Ravi. 2020. Goemotions: A dataset of fine-grained emotions. In *ACL*, pages 4040–4054.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In NAACL-HLT (1), pages 4171–4186.
- Ian S. Fischer. 2020. The conditional entropy bottleneck. *Entropy*, 22(9):999.
- Ziv Goldfeld, Ewout van den Berg, Kristjan H. Greenewald, Igor Melnyk, Nam Nguyen, Brian Kingsbury, and Yury Polyanskiy. 2019. Estimating information flow in deep neural networks. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 2299–2308. PMLR.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *ICLR (Poster)*.

- Beliz Gunel, Jingfei Du, Alexis Conneau, and Veselin Stoyanov. 2021. Supervised contrastive learning for pre-trained language model fine-tuning. In *ICLR*.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noisecontrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, volume 9 of *JMLR Proceedings*, pages 297–304. JMLR.org.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *SemEval@NAACL-HLT*, pages 39– 50.
- Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M. Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR (Poster)*.
- R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *ICLR*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735– 1780.
- Dou Hu, Yinan Bao, Lingwei Wei, Wei Zhou, and Songlin Hu. 2023. Supervised adversarial contrastive learning for emotion recognition in conversations. In *ACL* (1), pages 10835–10852. Association for Computational Linguistics.
- Dou Hu, Xiaolong Hou, Xiyang Du, Mengyuan Zhou, Lianxin Jiang, Yang Mo, and Xiaofeng Shi. 2022. Varmae: Pre-training of variational masked autoencoder for domain-adaptive language understanding. In *EMNLP (Findings)*, pages 6276–6286.
- Dou Hu, Lingwei Wei, Yaxin Liu, Wei Zhou, and Songlin Hu. 2024. Structured probabilistic coding. In *AAAI*, pages 12491–12501. AAAI Press.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomás Mikolov. 2017. Bag of tricks for efficient text classification. In EACL (2), pages 427–431. Association for Computational Linguistics.
- Kenji Kawaguchi, Zhun Deng, Xu Ji, and Jiaoyang Huang. 2023. How does information bottleneck help deep learning? In *ICML*, volume 202 of *Proceedings* of Machine Learning Research, pages 16049–16096. PMLR.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. In *NeurIPS*.
- Tae Kyun Kim. 2015. T test as a parametric statistic. *Korean journal of anesthesiology*, 68(6):540–546.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR (Poster)*.
- Diederik P. Kingma and Max Welling. 2014. Autoencoding variational bayes. In *ICLR*.
- Artemy Kolchinsky, Brendan D. Tracey, and David H. Wolpert. 2019. Nonlinear information bottleneck. *Entropy*, 21(12):1181.
- Lingpeng Kong, Cyprien de Masson d'Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. A mutual information maximization perspective of language representation learning. In *ICLR*.
- Yann LeCun, Yoshua Bengio, and Geoffrey E. Hinton. 2015. Deep learning. *Nat.*, 521(7553):436–444.
- Ralph Linsker. 1988. Self-organization in a perceptual network. *Computer*, 21(3):105–117.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2021. Variational information bottleneck for effective low-resource fine-tuning. In *ICLR*.
- Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. 2017. Adversarial training methods for semisupervised text classification. In *ICLR (Poster)*.
- Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 task 1: Affect in tweets. In SemEval@NAACL-HLT, pages 1–17.
- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiao-Dan Zhu, and Colin Cherry. 2016. Semeval-2016 task 6: Detecting stance in tweets. In SemEval@NAACL-HLT, pages 31–41.
- Xue Bin Peng, Angjoo Kanazawa, Sam Toyer, Pieter Abbeel, and Sergey Levine. 2019. Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow. In *ICLR (Poster)*. OpenReview.net.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Lukasz Kaiser, and Geoffrey E. Hinton. 2017. Regularizing neural networks by penalizing confident output distributions. In *ICLR (Workshop)*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In ACL, pages 527–536.
- Ruggero Ragonesi, Riccardo Volpi, Jacopo Cavazza, and Vittorio Murino. 2021. Learning unbiased representations via mutual information backpropagation. In *CVPR Workshops*, pages 2729–2738.

- Alexandre Ramé and Matthieu Cord. 2021. DICE: diversity in deep ensembles via conditional redundancy adversarial estimation. In *ICLR*.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. Semeval-2017 task 4: Sentiment analysis in twitter. In *SemEval@ACL*, pages 502–518. Association for Computational Linguistics.
- Michael Roth, Talita Anthonio, and Anna Sauer. 2022. Semeval-2022 task 7: Identifying plausible clarifications of implicit and underspecified phrases in instructional texts. In SemEval@NAACL, pages 1039–1049.
- Klaus R Scherer and Harald G Wallbott. 1994. Evidence for universality and cultural variation of differential emotion response patterning. *Journal of personality and social psychology*, 66(2):310.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 2256–2265. JMLR.org.
- Thomas Steinke and Lydia Zakynthinou. 2020. Reasoning about generalization via conditional mutual information. In *COLT*, volume 125 of *Proceedings of Machine Learning Research*, pages 3437–3452. PMLR.
- Matteo Terzi, Alessandro Achille, Marco Maggipinto, and Gian Antonio Susto. 2021. Adversarial training reduces information and improves transferability. In *AAAI*, pages 2674–2682. AAAI Press.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. 2020. What makes for good views for contrastive learning? In *NeurIPS*.
- Naftali Tishby, Fernando C Pereira, and William Bialek. 1999. The information bottleneck method. In *Proc. of the 37th Allerton Conference on Communication and Computation*.
- Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. In 2015 *ieee information theory workshop (itw)*, pages 1–5.
- Michael Tschannen, Josip Djolonga, Paul K. Rubenstein, Sylvain Gelly, and Mario Lucic. 2020. On mutual information maximization for representation learning. In *ICLR*. OpenReview.net.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *CoRR*, abs/1807.03748.
- Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 9929–9939. PMLR.

- Aolin Xu and Maxim Raginsky. 2017. Informationtheoretic analysis of generalization capability of learning algorithms. *Advances in Neural Information Processing Systems*, 30.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). In *SemEval@NAACL-HLT*, pages 75–86.

Appendix Overview

In the appendix, we will 1) briefly illustrate related work, 2) provide detailed experimental setups, and 3) report detailed experimental results.

A Related Work

Information-theoretic representation learning has become the workhorse of several impressive deep learning achievements, ranging from practical applications (Pereyra et al., 2017; Alemi et al., 2017; Belghazi et al., 2018) to theoretical investigations (Xu and Raginsky, 2017; Steinke and Zakynthinou, 2020; Kawaguchi et al., 2023). The aim is typically to discover useful and informative latent representations in a principled and systematic manner. In the context of the training principle, informationtheoretic representation learning approaches can be broadly categorized into InfoMax-based and IBbased methods.

In the field of machine learning, the information maximization (InfoMax) principle was first introduced by Linsker (1988). It suggests that a linear or nonlinear network can be viewed as an information channel and the goal is to transmit as much information as possible from the input data through the network. In recent years, InfoMax has extensive applications in the field of self-supervised representation learning (Chen et al., 2016; van den Oord et al., 2018; Hjelm et al., 2019; Tschannen et al., 2020; Kong et al., 2020). Chen et al. (2016) extend generative adversarial networks to learn disentangled representations by similarly using Info-Max. Tschannen et al. (2020) study the usage of InfoMax in representation learning and shows the success depend on the inductive bias by the model architectures and mutual information (MI) estimators. Kong et al. (2020) explore several embedding models based on InfoMax, and introduce an objective that maximizes MI between the sentence representation and the n-grams. In supervised scenarios, minimizing the standard cross-entropy is actually equivalent to maximizing the mutual information between the representations and the target task (Achille and Soatto, 2018; Kolchinsky et al., 2019; Boudiaf et al., 2020). Moreover, Pereyra et al. (2017) add the negative entropy to the negative log-likelihood to penalize overconfident output distributions in a supervised manner.

The information bottleneck (IB) principle (Tishby et al., 1999; Tishby and Zaslavsky, 2015) aims to balance the trade-off between the compres-

sion of representation and the power of predicting. Alemi et al. (2017) propose an efficient variational estimation method VIB based on the IB principle. Peng et al. (2019) use IB to constrain information flow in the discriminator for better adversarial learning. Then, Mahabadi et al. (2021) use VIB to suppress irrelevant features when fine-tuning on low-resource language understanding tasks. For tractable application of IB in a continuous setting, Belghazi et al. (2018) propose a mutual information neural estimation method to estimate MI and design an IB principle based on the estimator. Ragonesi et al. (2021) employ MINE-IB to learn unbiased representations. Fischer (2020) and Ramé and Cord (2021) introduce a conditional mutual information term to alleviate the over- and under-compression issue of IB. An et al. (2023) explicitly use the conditional entropy of the stochastic embedding as a confidence indicator and encourage the model to assign larger variance to more certain inputs. Recently, Hu et al. (2024) propose a probabilistic coding framework to mitigate the over-compression issue of IB by simultaneously performing compression and prediction using a shared learnable network. They also introduced a regularization term to promote class-level uniformity for better task prediction. Moreover, variational autoencoder (VAE) (Kingma and Welling, 2014) can be seen as a special case of an unsupervised VIB and have shown impressive results in self-supervised learning (Sohl-Dickstein et al., 2015; Higgins et al., 2017; Hu et al., 2022).

B Detailed Experimental Setups

B.1 Downstream Datasets and Tasks

We conduct experiments on 13 natural language understanding benchmarks, including 10 classification benchmarks (i.e., emoji prediction, hate speech detection, irony detection, offensive language detection, sentiment analysis, stance detection, and 4 categorical emotion analysis tasks from different domains), and 3 regression benchmarks (i.e., semantic similarity prediction, plausible clarifications ranking, and dimensional emotion analysis). The detailed statistics are shown in Table 6.

The descriptions of 10 classification benchmarks are listed as follows: *EmojiEval* (Barbieri et al., 2018) is designed for emoji prediction, which aims to predict its most likely emoji given a tweet. *EmotionEval* (Mohammad et al., 2018) involves detecting the emotion evoked by a tweet and is based on

Dataset	Task	# Label	# Train	# Val	# Test	# Total
Classification	Benchmarks					
EmojiEval	Emoji prediction	20	45,000	5,000	50,000	100,000
EmotionEval	Social emotion detection	4	3,257	374	1,421	5,052
HatEval	Hate speech detection	2	9,000	1,000	2,970	12,970
IronyEval	Irony detection	2	2,862	955	784	4,601
OffensEval	Offensive language detection	2	11,916	1,324	860	14,100
SentiEval	Sentiment analysis	3	45,389	2,000	11,906	59.295
StanceEval	Stance detection	3	2,620	294	1,249	4,163
ISEAR	Emotion reaction prediction	7	3,066	767	3,833	7,666
MELD	Conversational emotion recognition	7	9,989	1,109	2,610	13,708
GoEmotions	Fine-grained emotion detection	28	36,308	4,548	4,591	45,447
Regression Be	nchmarks					
STS-B	Semantic similarity prediction	-	7,000	1,500	1,400	9,900
CLAIRE	Plausible clarification ranking	-	19,975	2,500	2,500	24,975
EmoBank	Dimensional Emotion Analysis	-	8,062	1,000	1,000	10062

Table 6: The statistics of all datasets and tasks.

the Affects in Tweets conducted during SemEval-2018. Following Barbieri et al. (2020), the most common four emotions (i.e., anger, joy, sadness, and optimism) are selected as the label sets. HatEval (Basile et al., 2019) stems from SemEval-2019 HatEval and is used to predict whether a tweet is hateful towards immigrants or women. IronyEval (Hee et al., 2018) is from SemEval-2018 Irony Detection and consists of identifying whether a tweet includes ironic intents or not. OffensEval (Zampieri et al., 2019) is from SemEval-2019 OffensEval and involves predicting if a tweet contains any form of offensive language. SentiEval (Rosenthal et al., 2017) comes from SemEval-2017 and includes data from previous runs (2013, 2014, 2015, and 2016) of the same SemEval task. The goal is to determine if a tweet is positive, negative, or neutral. StanceEval (Mohammad et al., 2016) involves determining if the author's text has a favorable, neutral, or negative position towards a proposition or target. ISEAR (Scherer and Wallbott, 1994) is from International Survey On Emotion Antecedents And Reactions project and contains reports on seven emotions each by close to 3000 respondents in 37 countries on all 5 continents. It aims to predict the emotion reaction. Due to the lack of a predefined split in the original corpus, we use the processed version by Hu et al. (2024) that randomly splits the dataset into train/valid/test sets in a ratio of 4:1:5 based on the label distribution. MELD (Poria et al., 2019) contains multi-party conversation videos collected from Friends TV series. It is used to detect emotions in each utterance. The corpus contains many types of context, including dialogue, speaker, and multi-modal signals. Following Hu et al. (2024), this paper only considers the contextfree textual utterance to better evaluate sentence

classification performance. *GoEmotions* (Demszky et al., 2020) is a corpus of comments from Reddit, with human annotations to 27 emotion categories or neutral. It is used fine-grained emotion prediction. Following Hu et al. (2024), nearly 16% of multilabel data was removed from the source corpus to better evaluate the performance of multi-class classification.

The descriptions of 3 regression benchmarks are listed as follows: STS-B (Cer et al., 2017) is a collection of English sentence pairs drawn from news headlines, video and image captions, and natural language inference data. The semantic similarity prediction task is to predict the semantic textual similarity score from 0 (very dissimilar) to 5 (very similar) given each sentence pair. CLAIRE (Roth et al., 2022) dataset consists of manually clarified how-to guides from wikiHow⁵ with generated alternative clarifications and human plausibility judgements. The goal of plausible clarifications ranking task is to predict the continuous plausibility score on a scale from 1 (very implausible) to 5 (very plausible) given the clarification and its context. In our experiments, a special token pair (i.e., <e>and </e>) is introduced as the boundary of filler words. Emobank (Buechel and Hahn, 2017) is a large-scale text corpus across 6 domains and 2 perspectives and manually annotated with continuous VAD scores. Each sentence has three scores representing VAD in the range of 1 to 5.

B.2 Comparison Methods

We compare against the 4 universal models (i.e., SVM, FastText, BiLSTM, and GPT-3.5) and 7 representative deep representation learning tech-

⁵https://www.wikihow.com/

	Hyperparameter	EmojiEval	EmotionEval	HatEval	IronyEval	OffensEval	SentiEval	StanceEval	ISEAR	MELD	GoEmotions	STS-B	CLAIRE	EmoBank
T	Trade-off weight β	1	0.1	10	0.01	0.01	0.1	0.1	0.1	0.1	0.1	-	-	-
	Temperature τ	0.1	0.1	0.1	1	0.1	0.5	0.1	0.1	0.1	0.1	-	-	-
ER	Perturbation rate	1	1	1	1	0.1	1	1	1	1	1	-	-	-
щ	Perturbation radius ϵ	0.1	1	0.1	5	5	1	1	5	1	0.1	-	-	-
	Weight decay	0	0.001	0.01	0.001	0	0	0.001	0	0.001	0	-	-	-
-	Trade-off weight β	0.01	1	10	1	1	0.01	0.1	0.1	1	0.1	0.001	0.01	0.01
R	Temperature τ	0.1	0.5	0.1	1	1	0.1	0.5	0.1	1	0.1	0.1	1	1
BE	Perturbation rate	1	0.1	1	0.1	1	1	1	1	1	1	1	0.1	0.1
Ro]	Perturbation radius ϵ	0.1	5	0.1	0.1	1	1	0.1	0.1	1	1	5	0.1	1
	Weight decay	0	0	0.01	0.01	0.001	0	0	0	0	0	0	0	0

Table 7: Hyperparameters of CIFM on 13 natural language understanding tasks.

Methods	EmojiEval	EmotionEval	HatEval	IronyEval	OffensEval	SentiEval	StanceEval	ISEAR	MELD	GoEmotions	Avg.
BERT backbone											
CIFM	24.28±0.39	77.74±1.07	59.22±1.75	65.51±1.61	80.67±0.99	70.66±0.28	67.81±0.97	71.51 ± 0.51	41.72 ± 0.80	48.42 ± 0.74	60.75
- w/o CIM	24.05±0.35	77.30 ± 0.62	$55.69{\scriptstyle\pm2.11}$	60.89 ± 3.11	80.35 ± 1.19	$70.12{\scriptstyle \pm 0.73}$	65.66 ± 1.12	$67.65{\scriptstyle \pm 0.76}$	$40.30{\scriptstyle \pm 0.74}$	47.57 ± 0.66	58.96
- w/o CIM & IFM	22.17 ± 0.88	75.13±1.73	$46.96{\scriptstyle \pm 5.69}$	59.04±4.55	$80.26{\scriptstyle \pm 1.26}$	$70.57{\scriptstyle \pm 0.79}$	65.38 ± 0.57	$66.96{\scriptstyle \pm 0.64}$	$40.13{\scriptstyle \pm 0.49}$	46.11 ± 0.72	57.27
RoBERTa backbone											
CIFM	32.32±0.87	79.63±0.57	$59.55{\scriptstyle\pm0.91}$	63.02±3.81	80.96±0.55	72.93 ± 0.27	69.01 ± 0.80	$71.66{\scriptstyle \pm 0.68}$	43.99 ± 1.10	49.51±0.31	62.26
- w/o CIM	30.90±1.27	78.72 ± 0.72	$58.84{\scriptstyle\pm2.05}$	$61.31{\scriptstyle\pm3.60}$	$78.43{\scriptstyle \pm 0.50}$	$72.51{\scriptstyle \pm 0.72}$	67.10 ± 1.01	$70.59{\scriptstyle \pm 0.38}$	$43.23{\scriptstyle\pm1.61}$	48.00 ± 1.10	60.96
- w/o CIM & IFM	30.67±1.48	77.37 ± 1.04	$43.88{\scriptstyle\pm4.57}$	47.01±7.05	$80.29{\scriptstyle \pm 0.35}$	$72.74{\scriptstyle\pm0.40}$	$67.60{\scriptstyle \pm 0.74}$	$70.42{\scriptstyle \pm 0.47}$	$41.30{\scriptstyle \pm 1.67}$	$47.96{\scriptstyle \pm 0.73}$	57.92

Table 8: Ablation results (%) on 10 classification benchmarks. We bolded best results for each task under BERT and RoBERTa backbones.

Mathada	ST	S-B	CLA	IRE		A		
wiethous	Spearman Pearson		Spearman Pearson		Pearson (v)	Pearson (a)	Pearson (d)	Avg.
CIFM	88.94±0.38	89.52±0.32	53.62±0.51	52.88±0.30	81.64±0.84	57.06±0.95	51.60±1.11	68.64
- w/o CIM	88.37±0.43	89.03 ± 0.29	51.71 ± 0.80	50.36±1.52	79.33 ± 1.90	55.01±1.82	47.35±3.67	66.77
- w/o CIM & IFM	$88.44{\scriptstyle \pm 0.42}$	$89.10{\scriptstyle \pm 0.37}$	$51.45{\scriptstyle \pm 1.00}$	$50.82{\scriptstyle \pm 1.36}$	$80.37{\scriptstyle\pm1.58}$	$55.83{\scriptstyle \pm 0.68}$	$49.96{\scriptstyle \pm 3.65}$	67.32

Table 9: Ablation results (%) on 3 regression benchmarks with RoBERTa backbone.

	Methods		SentiEval \rightarrow Others								
Cls.	Pre-train Obj.	EmojiEval	EmotionEval	HatEval	IronyEval	OffensEval	StanceEval	ISEAR	MELD	GoEmotions	Avg.
	CE	4.83	38.15	58.86	52.60	76.60	37.37	36.11	25.03	11.38	37.88
Linear	CE+AT	4.89	37.81	57.54	51.62	78.14	37.26	37.96	26.47	12.07	38.20
	VIB	3.81	42.14	57.38	50.51	76.08	36.69	34.54	23.90	8.55	37.07
	MEIB	3.51	37.33	59.57	47.42	75.33	36.44	30.50	21.21	6.03	35.26
	CIFM	8.69	43.04	56.42	53.75	77.88	39.59	40.60	28.98	15.02	40.44
	CE	24.37	73.83	51.49	64.82	80.07	62.67	60.52	41.71	36.66	55.13
	CE+AT	25.61	74.21	49.09	66.62	80.43	62.70	61.17	41.34	37.41	55.40
CNN	VIB	16.93	72.09	54.66	64.11	79.51	59.69	57.22	40.51	32.95	53.07
	MEIB	13.94	70.81	54.32	62.25	78.78	57.44	52.48	38.54	29.87	50.94
	CIFM	26.43	74.88	50.12	65.02	80.64	63.57	62.81	43.20	39.23	56.21

Table 10: Transferability results (%) of *SentiEval* \rightarrow *Others* with RoBERTa backbone. Pre-train Obj. indicates the pre-training objective. Cls. refers to the classifier when fine-tuning. Best results for linear (i.e., Linear) and nonlinear (i.e., CNN) classifiers are highlighted in bold.

nologies (i.e., CE/MSE, CE+CP, CE/MSE+AT, CE+SCL, VIB, MINE-IB, and MEIB) with 2 different backbone models. VIB, MINE-IB, and MEIB belong to IB-based methods. The descriptions of these methods are listed as follows:

SVM (Cortes and Vapnik, 1995) is a machine learning algorithm with a hinge loss that aims to find the best hyperplane to separate data points into different classes. **FastText** (Joulin et al., 2017) is an efficient classification method with negative log-likelihood loss based on n-gram features and a hierarchical softmax. **BiLSTM** is a bidirectional recurrent neural network (Hochreiter and Schmidhuber, 1997). **GPT-3.5**⁶ is an enhanced generative pre-trained transformer model based on textdavinci-003, optimized for chatting.

CE/MSE means a fine-tuned baseline with a cross-entropy (CE) loss for classification or a mean squared error (MSE) loss for regression. **CE+CP** (Pereyra et al., 2017) is an entropy regularization method that fits a deterministic network by optimizing an objective that combines the CE loss with a confidence penalty term. **MSE/CE+AT** (Miyato et al., 2017) uses CE/MSE with classical adversarial training. **CE+SCL** (Gunel et al., 2021) combines CE and supervised contrastive learning (SCL) (Khosla et al., 2020). SCL allows for multiple positives per anchor, thus adapting contrastive learning to the fully supervised setting. **VIB** (Alemi et al., 2017)

⁶https://chat.openai.com

2017; Mahabadi et al., 2021) is an efficient variational estimation method based on IB (Tishby and Zaslavsky, 2015). **MINE-IB** (Belghazi et al., 2018) is a neural estimation method based on IB with a continuous setting. **MEIB** (An et al., 2023) is a variational approach to stochastic embedding in which maximum entropy acts as the bottleneck.

B.3 Implementation Details

All methods are conducted with the epoch number of 20, total batch size of 128, and maximum token length of 128. The maximum patience of early stopping is set to 5 epochs. The dropout rate is set to 0.2. For the default InfoNCE estimator used in IFM, the trade-off parameter β is searched from {0.01, 0.1, 1, 10} for classification and {0.001, 0.01, 0.1} for regression, and the temperature τ is searched from {0.1, 0.5, 1}. For the adversarial estimator used in CIM, the L_q norm constraint is L_2 , the perturbation rate is searched from {0.1, 1}, and the perturbation radius is searched from {0.1, 1, 5}. Table 7 shows the best parameters of our method with RoBERTa and BERT backbones.

For the compared GPT-3.5, we present the zeroshot results of the GPT-3.5-turbo snapshot from June 13th 2023 based on specific inputs, including task descriptions, instructions, and evaluation texts.

C Supplementary Results

C.1 Details of Ablation Studies

Table 8 and Table 9 report detailed ablation results on classification and regression tasks, respectively.

C.2 Details of Transferability Evaluation

Table 10 shows detailed transferability results of $SentiEval \rightarrow Others$.