# **TCGM: An Information-Theoretic Framework for** Semi-Supervised Multi-Modality Learning

Xinwei Sun<sup>\*1</sup>, Yilun Xu<sup>\*2</sup>, Peng Cao<sup>2</sup>, Yuqing Kong  $(\boxtimes)^2$ , Lingjing Hu<sup>3</sup>, Shanghang Zhang  $(\boxtimes)^4$ , and Yizhou Wang<sup>2,5</sup>

<sup>1</sup> Microsoft Research-Asia xinsun@microsoft.com Center on Frontiers of Computing Studies, Adv. Inst. of Info. Tech, Dept. of Computer Science, Peking University {xuyilun,caopeng2016, yuqing.kong, Yizhou.Wang}@pku.edu.cn <sup>3</sup> Yanjing Medical College, Capital Medical University hulj@ccmu.edu.cn <sup>4</sup> UC Berkeley shz@eecs.berkeley.edu <sup>5</sup> Deepwise AI Lab

Abstract. Fusing data from multiple modalities provides more information to train machine learning systems. However, it is prohibitively expensive and timeconsuming to label each modality with a large amount of data, which leads to a crucial problem of semi-supervised multi-modal learning. Existing methods suffer from either ineffective fusion across modalities or lack of theoretical guarantees under proper assumptions. In this paper, we propose a novel information-theoretic approach - namely, Total Correlation Gain Maximization (TCGM) - for semisupervised multi-modal learning, which is endowed with promising properties: (i) it can utilize effectively the information across different modalities of unlabeled data points to facilitate training classifiers of each modality (ii) it has theoretical guarantee to identify Bayesian classifiers, i.e., the ground truth posteriors of all modalities. Specifically, by maximizing TC-induced loss (namely TC gain) over classifiers of all modalities, these classifiers can cooperatively discover the equivalent class of ground-truth classifiers; and identify the unique ones by leveraging limited percentage of labeled data. We apply our method to various tasks and achieve state-of-the-art results, including the news classification (Newsgroup dataset), emotion recognition (IEMOCAP and MOSI datasets), and disease prediction (Alzheimer's Disease Neuroimaging Initiative dataset).

Keywords: Total Correlation, Semi-supervised, Multi-modality, Conditional Independence, Information intersection

#### 1 Introduction

Learning with data from multiple modalities has the advantage to facilitate information fusion from different perspectives and induce more robust models, compared with only using a single modality. For example, as shown in Figure 1, to diagnose whether a patient has a certain disease or not, we can consult to its X-ray images, look into its medical

<sup>\*</sup> Equal Contribution.

records, or get results from clinical pathology. However, in many real applications, especially in some difficult ones (e.g. medical diagnosis), annotating such large-scale training data is prohibitively expensive and time-consuming. As a consequence, each modality of data may only contain a small proportion of labeled data from professional annotators, leaving a large proportion of unlabeled. This leads to an essential and challenging problem of semi-supervised multi-modality learning: *how to effectively train accurate classifiers by aggregating unlabeled data of all modalities?* 

To achieve this goal, many methods have been proposed in the literature, which can be roughly categorized into two branches: (i) co-training strategy [6]; and (ii) learning joint representation across modalities in an unsupervised way [25,27]. These methods suffer from either too strong assumptions or loss of information during fusing. Specifically, the co-training strategy relies largely on the "compatible" assumption that the conditional distributions of the data point labels in each modality are the same, which may not be satisfied in the real settings, as self-claimed in [6]; while the latter branch of methods fails to capture the higher-order dependency among modalities, hence may end up in learning a trivial solution that maps all the data points to the same representation.



**Fig. 1.** Multiple modalities are independent conditioning on the ground truth; Ground truth is the "information intersection" of all of the modalities.

A common belief in multi-modality learning [22,6,13,20] is that conditioning on ground truth label *Y*, these modalities are conditionally independent, as illustrated in Figure 1. For example, to diagnose if one suffers from a certain disease, an efficient way is to leverage as many as modalities that are related to the disease, e.g., X-ray image, medical records and the clinical pathology. Since each modality captures the characteristics of the disease from different aspects, the information extracted from these modalities, in addition to the label, are not necessarily correlated with each other. This suggests that the ground truth label can be regarded as the "information intersection" across all the modalities, i.e., the amount of agreement shared by all the modalities.

Inspired by such an assumption and the fact that the *Total Correlation* [28] can measure the amount of information shared by M ( $M \ge 2$ ) variables, in this paper, we propose Total Correlation Gain (TCG), which is a function of classifiers of all the modalities, as a surrogate goal for maximization of mutual information, in order to infer the ground-truth labels (i.e., information intersection among these modalities). Based on the proposed TCG, we devise an information-theoretic framework called *Total Correlation Gain Maximization* (TCGM) for semi-supervised multi-modal learning. By maximizing TCG among all the modalities, the classifiers for different modalities. It can be

proved that the optimal classifiers for such a Total Correlation Gain are equivalent to the Bayesian posterior classifiers given each modality under some permutation function. With further leverage of labeled data, we can identify the Bayesian posterior classifiers. Furthermore, we devise an aggregator that employs all the modalities to forecast the labels of data. A simulated experiment is conducted to verify this theoretical result.

We apply TCGM on various tasks: (i) News classification with three pre-processing steps as different modalities, (ii) Emotion recognition with videos, audios, and texts as three modalities and (iii) disease prediction on medical imaging with the Structural magnetic resonance imaging (sMRI) and Positron emission tomography (PET) modalities. On these tasks, our method consistently outperforms the baseline methods especially when a limited percentage of labeled data are provided. To validate the benefit of jointly learning, we visualize that some cases of Alzheimer's Disease whose label are difficult to be predicted via supervised learning with single modality; while our jointly learned single modal classifier is able to correctly classify such hard samples.

The contributions can be summarized as follows: (i) We propose a novel informationtheoretic approach TCGM for semi-supervised multi-modality learning, which can effectively utilize information across all modalities. By maximizing the total correlation gain among all the modalities, the classifiers for different modalities cooperatively discover the information intersection across all the modalities - the ground truth. (ii) To the best of our knowledge, **TCGM** is the first in the literature that can be theoretically proved that, under the conditional independence assumption, it can identify the groundtruth Bayesian classifier given each modality. Further, by aggregating these classifiers, our method can learn the Bayesian classifier given all modalities. (iii) We achieve the state-of-the-art results on various semi-supervised multi-modality tasks including news classification, emotion recognition and disease prediction of medical imaging.

# 2 Related Work

*Semi-supervised multi-modal learning* It is commonly believed in the literature that information of label is shared across all modalities. Existing work, which can be roughly categorized into two branches, suffers from either stronger but not reasonable assumptions or failure to capture the information (i.e., label) shared by all modalities. The first branch applies the co-training algorithm proposed by Blum et. al [6]. [17,3,12,21,11] use weak classifiers trained by the labeled data from each modality to bootstrap each other by generating labels for the unlabeled data. However, the underlying compatible condition of such a method, which assumes the same conditional distributions for data point labels in each modality, may not be consistent with the real settings.

The second branch of work [25,27,9,31,10,18] centers on learning joint representations that project unimodal representations all together into a multi-modal space in an unsupervised way and then using the labeled data from each modality to train a classifier to predict the label of the learned joint representation. A representative of such a framework is the soft-Hirschfeld-Gebelein-Rényi (HGR) framework [31], which proposed to maximize the correlation among non-linear representations of each modality. However, HGR only measures the linear dependence between pair modalities, since it follows the principle of maximizing the correlation between features of different modalities. In contrast, our framework, i.e., Total Correlation Gain Maximization can pursue information about higher-order dependence. Due to the above reasons, both branches can not avoid learning a naive solution that classifies all data points into the same class.

To overcome these limitations, we propose an information-theoretic loss function based on Total Correlation which can not only require the assumption in the first branch of work but also can be able to identify the ground-truth label, which is the information intersection among these modalities. Therefore, our method can avoid the trivial solution and can learn the optimal, i.e., the Bayesian Posterior classifiers of each modality.

Total Correlation/Mutual information maximization Total Correlation [28], as an extension of Mutual Information, measures the amount of information shared by M ( $M \ge 2$ ) variables. There are several works in the literature that have combined Mutual Information (M = 2) with deep learning algorithms and have shown superior performance on various tasks. Belghazi et al. [4] presents a mutual information neural estimator, which is utilized in a handful of applications based on the mutual information maximization (e.g., unsupervised learning of representations [14], learning node representations within graph-structured data [30]). Kong and Schoenebeck [19] provide another mutual information estimator in the co-training framework for the peer prediction mechanism, which has been combined with deep neural networks for crowdsourcing [8]. Xu et.al [32] proposes an alternative definition of information, which is more effective for structure learning. However, those three estimators can only be applied to two-view settings. To the best of our knowledge, there are no similar studies that focus on a general number of modalities, which is very often in real applications. In this paper, we propose to leverage Total Correlation to fill in such a gap.

# **3** Preliminaries

Notations Given a random variable X,  $\mathcal{X}$  denotes its realization space and  $x \in \mathcal{X}$  denotes an instance. The  $P_{\mathcal{X}}$  denotes the probability distribution function over  $\mathcal{X}$  and  $p(x) := dP_{\mathcal{X}}(x)$  denotes the density function w.r.t the Lebesgue measure. Further, given a finite set  $\mathcal{X}$ ,  $\Delta_{\mathcal{X}}$  denotes the set of all distributions over  $\mathcal{X}$ . For every integer M, [M] denotes the set  $\{1, 2, \ldots, M\}$ . For a vector  $\mathbf{v}$ ,  $v_i$  denotes its *i*-th element.

*Total Correlation* The Total Correlation (TC), as an extension of mutual information, measures the "amount of information" shared by  $M (\geq 2)$  random variables:

$$TC(X^{1},...,X^{m}) = \sum_{i=1}^{M} H(X^{i}) - H(X^{1},...,X^{M}),$$
(1)

where H is the Shannon entropy. As defined, the TC degenerates to mutual information when M = 2. The  $\sum_{i=1}^{M} H(X^i)$  measures the total amount of information when treating  $X^1, ..., X^M$  independently; while the  $H(X^1, ..., X^M)$  measures the counterpart when treating these M variables as a whole. Therefore, the difference between them implies the redundant information, i.e., the information shared by these M variables.

Similar to mutual information, the TC is equivalent to the Kullback-Leibler (KL)divergence between  $P_{\mathcal{X}^1 \times ... \times \mathcal{X}^M}$  and product of marginal distribution  $\otimes_{m=1}^M P_{\mathcal{X}^m}$ :

$$\operatorname{TC}(X^{1},...,X^{M}) = \operatorname{D}_{\operatorname{KL}}\left(d\mathbb{P}_{\mathcal{X}^{1}\times...\times\mathcal{X}^{M}} \mid\mid d\otimes_{m=1}^{M}\mathbb{P}_{\mathcal{X}^{m}}\right) = \mathbb{E}_{\mathbb{P}_{\mathcal{X}^{1}\times...\times\mathcal{X}^{M}}}\log\frac{d\mathbb{P}_{\mathcal{X}^{1}\times...\times\mathcal{X}^{M}}}{d\otimes_{m=1}^{M}\mathbb{P}_{\mathcal{X}^{m}}}$$
(2)

where  $D_{KL}(\mathcal{P} \mid\mid \mathcal{Q}) = \mathbb{E}_{\mathcal{P}} \log \frac{d\mathcal{P}}{d\mathcal{Q}}$ . Intuitively, larger KL divergence between joint and marginal distribution indicates more dependence among these M variables. To better characterize such a property, we give a formal definition of "Point-wise Total Correlation" (PTC):

**Definition 1 (Point-wise Total Correlation).** Given M random variables  $X^1, ..., X^M$ , the Point-wise Total Correlation on  $(x^1, ..., x^m) \in \mathcal{X}^1 \times ... \times \mathcal{X}^M$ , i.e.,  $\operatorname{PTC}(x^1, ..., x^M)$  is defined as  $\operatorname{PTC}(x^1, ..., x^M) = \log \frac{p(x^1, ..., x^M)}{p(x^1) ... p(x^M)}$ . Further, the  $R(x^1, ..., x^M) := \frac{p(x^1, ..., x^M)}{p(x^1) ... p(x^M)}$  is denoted as the joint-margin ratio.

*Remark 1.* The Point-wise Total Correlation can be understood as the point-wise distance between joint distribution and the marginal distribution. In more details, as noted from [15], by applying first-order Taylor-expansion, we have  $\log \frac{p(x)}{q(x)} \approx \log 1 + \frac{p(x)-q(x)}{q(x)} = \frac{p(x)-q(x)}{q(x)}$ . Therefore, the expected value of  $PTC(\cdot)$  can well measure the amount of information shared among these variables, which will be shown later in detailed.

For simplicity, we denote  $p^{[M]}(x) := p(x^1, ..., x^M)$  and  $q^{[M]}(x) := \prod_{i=1}^M p(x^i)$ . According to dual representation in [26], we have the following lower bound for KL divergence between p and q, and hence TC.

Lemma 1 (Dual version of *f*-divergence [26]).

$$D_{\mathrm{KL}}\left(p^{[M]} \mid\mid q^{[M]}\right) \ge \sup_{g \in \mathcal{G}} \mathbb{E}_{x \sim p^{[M]}}[g(x)] - \mathbb{E}_{x \sim q^{[M]}}\left[e^{g(x)-1}\right]$$
(3)

where  $\mathcal{G}$  is the set of functions that maps  $\mathcal{X}^1 \times \mathcal{X}^2 \times \cdots \times \mathcal{X}^M$  to  $\mathbb{R}$ . The equality holds if and only if  $g(x^1, x^2, \dots, x^M) = 1 + \operatorname{PTC}(x^1, \dots, x^M)$ , and the supremum is  $D_{\mathrm{KL}}(p^{[M]} || q^{[M]}) = \mathbb{E}_{p^{[M]}}(\operatorname{PTC})$ .

The Lemma 1 is commonly utilized for estimation of Mutual information [4] or optimization as variational lower bound in the machine learning literature. Besides, it also informs that the PTC is the optimal function to describe the amount of information shared by these M variables. Indeed, such shared information is the information intersection among these variables, i.e., *conditional on such information, these M variables are independent of each other*. To quantitatively describe this, we first introduce the *conditional total correlation* (CTC). Similar to TC, CTC measures the amount of information shared by these M variables conditioning on some variable Z:

**Definition 2** (Conditional Total Correlation (CTC)). Given M + 1 random variables  $X^1, ..., X^M, Z$ , the Conditional Total Correlation (CTC $(X^1, ..., X^M | Z)$ ) is defined as  $CTC(X^1, ..., X^M | Z) = \sum_{i=1}^M H(X^i | Z) - H(X^1, ..., X^M | Z)$ .

# 4 Method

Problem statement In the semi-supervised muli-modal learning scenario, we have access to an unlabeled dataset  $\mathcal{D}_u = \{x_i^{[M]}\}_i$  and a labeled dataset  $\mathcal{D}_l = \{(x_i^{[M]}, y_i)\}_i$ . Each label  $y_i \in \mathcal{C}$ , where  $\mathcal{C}$  denotes the set of classes. Each datapoint  $x_i^{[M]} :=$  
$$\begin{split} &\{x_i^1, x_i^2, \dots, x_i^M | x_i^m \in \mathcal{X}^m \} \text{ consists of } M \text{ modalities, where } \mathcal{X}^m \text{ denotes the domain of the } m\text{-th modality. Datapoints and labels in } \mathcal{D}_l \text{ are i.i.d. samples drawn from the joint distribution } U_{X^{[M]},Y}(x^1, x^2, \dots, x^M, y) := \Pr(X^1 = x^1, X^2 = x^2, \dots, X^M = x^M, Y = y). \text{ Data points in } \mathcal{D}_u \text{ are i.i.d. samples drawn from joint distribution } U_{X^{[M]},Y}(x^1, x^2, \dots, x^M, y) := \Pr(X^1 = x^1, X^2 = x^2, \dots, X^M = \sum_{c \in \mathcal{C}} U_{X^{[M]},Y}(x^1, x^2, \dots, x^M, y = c). \text{ Denote the prior of the ground truth labels by } \mathbf{p}^* = (\Pr(Y = c))_c. \text{ Upon the labeled and unlabeled datasets, our goal is to train } M \text{ classifiers } h^{[M]} := \{h^1, h^2, \dots, h^M\} \text{ and an aggregator } \zeta \text{ such that } \forall m, h^m : \mathcal{X}^m \to \Delta_{\mathcal{C}} \text{ predicts the ground truth } y \text{ based on a } m\text{-th modality } x^m \text{ and } \zeta : \mathcal{X}^1 \times \mathcal{X}^2 \times \dots \times \mathcal{X}^M \to \Delta_{\mathcal{C}} \text{ predicts the ground truth } y \text{ based on all of the modalities } x^{[M]}. \end{split}$$

*Outline* We will first introduce the assumptions regarding the ground truth label Y and prior distribution on  $(X^1, ..., X^M, Y)$  in section 4.1. In section 4.2, we will present our method, i.e., maximize the total correlation gain on unlabeled dataset  $\mathcal{D}_u$ . Finally, we will introduce our algorithm for optimization in section 4.3.

#### 4.1 Assumptions for Identification of Y

In this section, we first introduce two basic assumptions to ensure that the ground-truth label can be identified. According to Proposition 2.1 in [1], the label Y can be viewed as the generating factor of data X. Such a result can be extended to multiple variables (please refer supplementary for details), which implies that Y is the common generating factor of  $X^1, ..., X^M$ . Motivated by this, it is natural to assume that the ground truth label Y is the "information intersection" among  $X^1, ..., X^M$ , i.e., all of the modalities are independent conditioning on the ground-truth:

Assumption 1 (Conditional Independence). Conditioning on  $Y, X^1, X^2, \ldots, X^M$  are independent, i.e.,  $\forall x^1, \ldots, x^M$ ,  $\Pr(X^{[M]} = x^{[M]}|Y = c) = \prod_m \Pr(X^m = x^m|Y = c)$ , for any  $c \in C$ .

On the basis of this assumption, one can immediately get the conditional total correlation gain  $CTC(X^1, ..., X^M | Y) = 0$ . In other words, conditioning on Y, there is no extra information shared by these M modalities, which is commonly assumed in the literature of semi-supervised learning [22,6,13,20]. However, the Y may not be the unique information intersection among these M modalities. Specifically, the following lemma establishes the rules for such information intersection to hold:

**Lemma 2.** Given assumption 1,  $R(x^1, ..., x^M)$  (Joint-marginal ratio definition 1) has  $R(x^1, ..., x^M) = \sum_{c \in \mathcal{C}} \frac{\prod_m \Pr(Y=c|X^m=x^m)}{(\Pr(Y=c))^{M-1}}$ . Further, the optimal g in lemma 1 satisfies  $g(x^1, ..., x^M) = 1 + \log \sum_{c \in \mathcal{C}} \frac{\prod_m \Pr(Y=c|x^m)}{(\Pr(Y=c))^{M-1}}$ .

In other words, in addition to  $\{\Pr(Y = c | X^m = x^m)\}_m, \mathbf{p}_c^*$ , there are other solutions  $\{a_{x^1}, ..., a_{x^M}\}, r$  with  $a_{x^i} \in \Delta_{\mathcal{C}}$  (for  $i \in \mathcal{C}$ ) and  $r \in \Delta_{\mathcal{C}}$  that can make the g optimal, as long as its joint-marginal ratio is equal to the ground-truth one:

$$\sum_{c \in \mathcal{C}} \frac{\prod_{m} a_{x^{m}}}{(r_{c}^{a})^{M-1}} = R(x^{1}, ..., x^{M})$$
(4)

To make  $\{\Pr(Y = c | X^m = x^m)\}_m$ ,  $\mathbf{p}_c^*$  identifiable w.r.t a trivial permutation, we make the following trivial assumption on  $\Pr(X^1, ..., X^M, Y)$ .

**Assumption 2** (Well-defined Prior). The solutions  $\{a_{x^1}, ..., a_{x^M}\}, r^a$  and  $\{b_{x^1}, ..., b_{x^M}\}, r^b$  for Eq. (4) are equivalent under the permutation  $\prod$  on  $C: a_{x^m} = \prod b_{x^m}, r^a = \prod r^b$ .

7

#### 4.2 Total Correlation Gain Maximization (TCGM)

Assumption 1 indicates that the label Y is the generating factor of all modalities, and assumption 2 further ensures its uniqueness under permutation. Our goal is to learn the ground-truth label Y which is the information intersection among M modalities. In this section, we propose a novel framework, namely **Total Correlation Gain Maximization** (**TCGM**) to capture such an information intersection, which is illustrated in Figure. 2. To the best of our knowledge, we are the first to theoretically prove the identification of ground truth classifiers on semi-supervised multi-modality data, by generalizing [19,8] that can only handle two views in multi-view scenario. The high-level spirit is designing TC-induced loss over classifiers of every modality. By maximizing such a loss, these classifiers can converge to Bayesian posterior, which is the optimal solution of TC as expectation of the loss. First, we introduce the basic building blocks for our method.

Classifiers  $h^{[M]}$  In order to leverage the powerful representation ability of deep neural network (DNN), each classifier  $h^m(x^m; \Theta^m)$  is modeled by a DNN with parameters  $\Theta^m$ . For each modality m, we denote the set of all such classifiers by  $H^m$  and  $H^{[M]} := \{H^1, H^2, \ldots, H^M\}$ .

Modality Classifiers-Aggregator  $\zeta$  Given M classifiers for each modality  $h^{[M]}$  and a distribution  $\mathbf{p} = (p_c)_c \in \Delta_{\mathcal{C}}$ , the aggregator  $\zeta$  which predicts the ground-truth label by aggregating classifiers of all modalities, is constructed by  $\zeta(x^{[M]}; h^{[M]}, \mathbf{p}) =$ Normalize  $\left(\left(\frac{\prod_m h^m(x_i^m)_c}{(p_c)^{M-1}}\right)_c\right)$ , where Normalize $(\mathbf{v}) := \frac{\mathbf{v}}{\sum_c v_c}$  for all  $\mathbf{v} \in \Delta_{\mathcal{C}}$ .

*Reward Function*  $\mathcal{R}$  We define a reward function  $\mathcal{R} : \overbrace{\Delta_{\mathcal{C}} \times ... \times \Delta_{\mathcal{C}}}^{M} \to \mathbb{R}$  that measures the "amount of agreement" among these classifiers. Note that the desired classifiers should satisfy Eq. (4).

Inspired by Lemma 1, we can take the empirical total correlation gain of N samples, i.e., the lower bound of Total Correlation as our maximization function. Specifically, given a reward function  $\mathcal{R}$ , the empirical total correlation with respect to classifiers  $h^{[M]}$ , a prior  $\mathbf{p} \in \Delta_{\mathcal{C}}$  measures the empirical "amount of agreement" for these M classifiers at the same sample  $(x_i^1, ..., x_i^M) \in \mathcal{D}_u$ , with additionally a punishment of them on different samples:  $x_{i_1}^1 \in \mathcal{X}^1, ..., x_{i_M}^M \in \mathcal{X}^M$  with  $i_1 \neq i_2 \neq ... \neq i_M$ :

$$TCg[\mathcal{R}](\{x_i^{[M]}\}_{i=1}^N; h^{[M]}, p) := \frac{1}{N} \sum_i \mathcal{R}(h^1(x_i^1), ..., h^M(x_i^M)) - \frac{1}{N!/(N-M)!} \sum_{i_1 \neq i_2 \neq \cdots \neq i_M} e^{\mathcal{R}(h^1(x_{i_1}^1), ..., h^M(x_{i_M}^M)) - (5)}$$

for simplicity we denote  $TCg[\mathcal{R}](\{x_i^{[M]}\}_{i=1}^N; h^{[M]}, p)$  as  $TCg^{(N)}$ . Intuitively, we expect our classifiers to be consistent on the same sample; on the other hand, to disagree on different samples to avoid learning a trivial solution that classifies all data points into the same class.

**Definition 3** (Bayesian posterior classifiers/aggregator). The  $h_*^{[M]}$  and  $\zeta_*$  are called Bayesian posterior classifiers and Bayesian posterior aggregator if they satisfy

$$\forall m, h_*^m(x^m)_c = \Pr(Y = c | X^m = x^m); \ \zeta_*(x^{[M]})_c = \Pr(Y = c | X^{[M]} = x^{[M]}).$$

Note that from Eq. (5) that our maximization goal, i.e.,  $TCg^{(N)}$  relies on the form of reward function  $\mathcal{R}$ . The following Lemma tells us the form of optimal reward function, with which we can finally give an explicit form of  $TCq^{(N)}$ .

**Lemma 3.** The  $\mathcal{R}_*$  that maximizes the expectation of  $TCg^{(N)}$  can be represented as the Point-wise Total Correlation function, which is the function of Bayesian classifiers and the prior of ground truth labels  $(\mathbf{p}_c^*)_c: \mathcal{R}_*(h^1(x^1), ..., h^M(x^M)) = 1 + \operatorname{PTC}(x^1, ..., x^M) = 1 + \log \sum_{c \in \mathcal{C}} \frac{\prod_m h_*^m(x^m)_c}{(\mathbf{p}_c^*)^{M-1}}.$ 

*Total Correlation Gain* Bring  $\mathcal{R}_*$  to Eq. (5), we have:

$$TCg(\{x_i^{[M]}\}_{i=1}^N; h^{[M]}, p) := 1 + \frac{1}{N} \sum_i \log \sum_{c \in \mathcal{C}} \frac{\prod_m h^m (x_i^m)_c}{(p_c)^{M-1}} - \frac{1}{N!/(N-M)!} \sum_{i_1 \neq i_2 \neq \dots \neq i_M} \sum_{c \in \mathcal{C}} \frac{\prod_m h^m (x_{i_m}^m)_c}{(p_c)^{M-1}}.$$
 (6)

As inspired by Lemma 3, we have that these Bayesian posterior classifiers are maximizers of the expected total correlation gain. Therefore, we can identify the equivalent class of Bayesian posteriors by minimizing  $-TCg^{(N)}$  on unlabeled dataset  $\mathcal{D}_u$ . By additionally minimize expected cross entropy (CE) loss on  $\mathcal{D}_l$ , we can identify the unique Bayesian classifiers since they are respectively the minimizers of CE loss.

**Theorem 3** (Main theorem). Define the expected total correlation gain  $eTCg(h^{[M]}, p)$ :

$$eTCg(h^1,...,h^M,p) := \mathbb{E}_{x_i^{[M]} \sim U_{X^{[M]}}} \left( TCg(x_i^{[M]};h^{[M]},p) \right)$$

Given the conditional independence assumption 1 and well-defined prior assumption 2, we have that the maximum value of eTCg is Total Correlation of M modalities, i.e.,  $TC(X^1, ..., X^M)$ . Besides,

 $\begin{array}{l} \textbf{Ground-truth} \rightarrow \textbf{Maximizer}(h_*^{[M]}, \mathbf{p}^*) \text{ is a maximizer of } eTCg(h^{[M]}, p). \text{ In other} \\ words, \forall h^{[M]} \in H^{[M]}, p \in \Delta_{\mathcal{C}}, eTCg(h_*^1, ..., h_*^M, \mathbf{p}^*) \geq eTCg(h^1, ..., h^M, p). \text{ The} \\ corresponding optimal aggregator is \zeta_*, i.e., \zeta_*(x^{[M]})_c = \Pr(Y = c | X^{[M]} = x^{[M]}). \\ \textbf{Maximizer} \rightarrow (\textbf{Permuted}) \text{ Ground-truth } \text{ If the prior is well defined, then for any} \\ \hline \end{array}$ 

maximizer of eTCg,  $(\tilde{h}^{[M]}, \tilde{\mathbf{p}})$ , there is a permutation  $\tilde{\prod} : \mathcal{C} \to \mathcal{C}$  such that:

$$\tilde{h}^m(x^m)_c = \mathcal{P}(Y = \prod(c) | X^m = x^m), \ \tilde{\mathbf{p}}_c = \mathcal{P}(Y = \prod(c))$$
(7)

The proof is in Appendix A. Note from our main theorem that by maximizing the eTCg, we can get the total correlation of M modalities, which is the ground-truth label Y, and also the equivalent class of Bayesian posterior classifier under permutation function. In order to identify the Bayesian posterior classifiers, we can further minimize cross-entropy loss on labeled data  $\mathcal{D}_l$  since the Bayesian posterior classifiers are the only minimizers of the expected cross-entropy loss. On the other hand, compared with only using  $\mathcal{D}_l$  to train classifiers, our method can leverage more information from  $\mathcal{D}_u$ , which can be shown in the experimental result later.

#### 4.3 Optimization

Since eTCg is intractable, we alternatively maximize the empirical total correlation gain, i.e.,  $TCg^{(N)}$  to learn the optimal classifiers. To identify the unique Bayesian posteriors, we should further utilize labeled dataset  $\mathcal{D}_l$  in a supervised way. Our whole optimization process is shown in Appendix, which adopts iteratively optimization strategy that is roughly contains two steps in each round: (i) We train the *M* classifiers using the classic cross entropy loss on the labeled dataset  $\mathcal{D}_l$  and (ii) using our information-theoretic loss function  $\mathcal{L}_{TC}$  on the unlabeled dataset  $\mathcal{D}_u$ . To learn the Bayesian posterior classifiers more accurately, the (ii) can help to learn the equivalent class of Bayesian Posterior Classifiers and (i) is to learn the correct and unique classifiers. As shown in Figure 2, by optimizing  $\mathcal{L}_{TC}^{(B)}$  (Eq. (5) with *B* denoting the number of samples in each batch), we reward the *M* classifiers for their agreements on the same data point and punish the *M* classifiers for their agreements on different data points.

Loss function  $\mathcal{L}_{CE}$  for labeled data We use the classic cross entropy loss for labeled data. Formally, for a batch of data points  $\{x_i^{[M]}\}_{i=1}^B$  drawn from labeled data  $\mathcal{D}_l$ , the cross entropy loss  $\mathcal{L}_{CE}$  for each classifier  $h^m$  is defined as  $\mathcal{L}_{CE}(\{(x_i^{[M]}, y_i)\}_{i=1}^B; h^m) := \frac{1}{B}\sum_i -\log(h^m(x_i^m)_{y_i})$ .

Loss function  $\mathcal{L}_{TC}^{(B)}$  for unlabeled data For a batch of data points  $\{x_i^{[M]}\}_{i=1}^B$  drawn from unlabeled data  $\mathcal{D}_u$ , our loss function  $\mathcal{L}_{TC}^{(B)} := -TCg^{(B)}$  that is defined in Eq. (6) with N replaced by number of batch size B. When N is large, we only sample a fixed number of samples from product of marginal distribution to estimate the second term in Eq. (6), which makes training more amenable.



**Fig. 2.** Empirical Total Correlation Gain  $TCg^{(N)}$ 

Prediction After optimization, we can get the classifiers  $\{h_m\}_m$ . The prior  $\mathbf{p}_c$  can be estimated from data, i.e.,  $\mathbf{p}_c = \frac{\sum_{i \in [N]} \mathbf{1}(y_i = c)}{N}$ . Then based on Eq. (7), we can get the aggregator classifier  $\zeta$  for prediction. Specifically, given a new sample  $\tilde{x}^{[M]}$ , the predicted label is  $\tilde{y} := \arg \max_c \zeta(\tilde{x}^{[M]})_c$ .

*Time Complexity* The overall time complexity for optimizing our TCGM is linear scale to the number of modalities, i.e.,  $\mathcal{O}(M)$ . Please refer to appendix for detailed analysis.

# **5** Preliminary experiments

We first conduct a simulated experiment on synthetic data to validate our theoretical result of TCGM. Specifically, we will show the effectiveness of Total Correlation Gain TCq for unsupervised clustering of data. Further, with few labeled data, our TCGM can give accurate classification. In more detail, we synthesize the data of three modalities from a specific Gaussian distribution  $P(X^i|y)$  (i = 1, 2, 3). The clustering accuracy is calculated as classification accuracy by assuming the label is known. As shown in Figure 3, our method TCGM has competitive performance compared to well established clustering algorithms Kmeans++ [2] and spectral clustering [24]. Based on the promising unsupervised learning result, as shown by the light blue line (the top line) in Figure 3, our method can accurately classify the data



Fig. 3. Clustering and classification accuracy.

even with only a small portion of labeled data. In contrast, **HGR** [31] degrades since it can only capture the linear dependence and fail when higher-order dependency exists.

# 6 Applications

In this section, we evaluate our method on various multi-modal classification tasks: (i) News classification (Newsgroup) (ii) Emotion recognition: **IEMOCAP**, **MOSI** and (iii) Disease prediction of Alzheimer's Disease on 3D medical Imaging: **Alzheimer's Disease Neuroimaging Initiative (ADNI)**. Our method **TCGM** is compared with : **CE** separately trains classifiers for each modality by minimizing cross entropy loss of only labeled data; **HGR** [31] learns representation by maximizing correlation of different modalities; and **LMF** [23] performs multimodal fusion using low-rank tensors. The optimal hyperparameters are selected according to validation accuracy, among which the learning rate is optimized from {0.1, 0.01, 0.001, 0.0001}. All experiments are repeated five times with different random seeds. The mean test accuracies and standard deviations of single classifiers the aggregator ( $\zeta$ ) are reported.

## 6.1 News Classification

*Dataset* **Newsgroup** [5] <sup>6</sup> is a group of news classification datasets. Following [16], each data point has three modalities, PAM, SMI and UMI, collected from three different preprocessing steps <sup>7</sup>. We evaluate **TCGM** and the baseline methods on 3 datasets from **Newsgroup**: News-M2, News-M5, News-M10. They contain 500, 500, 1000 data points with 2, 5, 10 categories respectively. Following [33], we use 60% for training, 20% for validation and 20% for testing for all of these three datasets.

<sup>&</sup>lt;sup>6</sup> http://qwone.com/~jason/20Newsgroups/

<sup>&</sup>lt;sup>7</sup> PAM (Partitioning Around Medoïds preprocessing), SMI (Supervised Mutual Information preprocessing) and UMI (Unsupervised Mutual Information preprocessing)

*Implementation details* We synthesize two different label rates (the percentage of labeled data points in each modality):  $\{10\%, 30\%\}$  for each dataset. We follow [33] for classifiers. Adam with default parameters and learning rate  $\gamma_u = 0.0001, \gamma_l = 0.01$  is used as the optimizer during training. Batch size is set to 32. We further compare with two additional baselines: **VAT** [29] uses adversarial training for semi-supervised learning; **PVCC** [33] that considers the consistency of data points under different modalities.

	PAM	SMI	LMI	All		PAM	SMI	LMI	All		PAM	SMI	LMI	All
CE	$78.20 \pm 4.21$	$84.20 \pm 7.01$	$77.40 \pm 5.03$	$92.00 \pm 2.24$	CE	$64.20 \pm 3.90$	$67.80 \pm 8.20$	$53.40 \pm 4.72$	$78.00 \pm 6.89$	CE	$27.80 \pm 4.40$	$31.70 \pm 4.86$	$31.00 \pm 3.74$	$38.10 \pm 5.47$
VAT	$80.60\pm6.02$	$84.00 \pm 4.74$	$75.80 \pm 2.77$	$90.80 \pm 2.05$	VAT	$65.60 \pm 1.82$	$66.80 \pm 8.38$	$52.40 \pm 3.29$	$79.40 \pm 5.94$	VAT	$30.70 \pm 1.25$	$30.60 \pm 4.02$	$27.40 \pm 3.05$	$37.20 \pm 5.20$
PVCC	$80.80 \pm 2.77$	$80.60 \pm 2.97$	$78.40 \pm 7.64$	$92.20\pm3.35$	PVCC	$67.80 \pm 8.38$	$67.00 \pm 1.22$	$54.60 \pm 6.02$	$79.00 \pm 6.28$	PVCC	$31.80 \pm 2.89$	$34.10\pm3.11$	$31.10 \pm 4.14$	$46.40 \pm 5.08$
HGR	-	-	-	$92.40 \pm 2.19$	HGR	-	-	-	$79.80 \pm 7.40$	HGR	-	-	-	$38.20 \pm 3.72$
TCGM	$93.60 \pm 1.14$	$93.80 \pm 2.17$	$93.00\pm1.58$	$98.20 \pm 1.10$	TCGM	$89.00\pm2.24$	$90.60\pm5.27$	$76.00\pm3.54$	$97.20 \pm 1.10$	TCGM	$45.50\pm6.29$	$47.60\pm3.73$	$52.70\pm5.12$	$64.10\pm2.16$
	News-M2, Label Rate 10%				News-M5, Label Rate 10%				News-M10, Label Rate 10%					
	PAM	SMI		4.19		Dill	ener	T MT			PAM	CMI	1 MI	A11
		3011	LMI	All		PAM	5011	LMI	All		174.04	3811	Lan	
CE	$86.40 \pm 2.30$	91.20 ± 1.92	85.40 ± 1.95	96.40 ± 1.82	CE	85.80 ± 3.63	81.80 ± 2.86	63.80 ± 5.40	96.20 ± 3.11	CE	49.00 ± 4.40	50.80 ± 2.49	51.20 ± 2.51	69.50 ± 4.57
CE VAT	86.40 ± 2.30 89.00 ± 3.87	91.20 ± 1.92 91.20 ± 3.56	85.40 ± 1.95 84.40 ± 2.70	96.40 ± 1.82 96.20 ± 1.92	CE VAT	85.80 ± 3.63 85.40 ± 2.88	81.80 ± 2.86 83.00 ± 2.55	63.80 ± 5.40 60.20 ± 8.98	96.20 ± 3.11 97.00 ± 2.83	CE VAT	49.00 ± 4.40 46.80 ± 4.31	50.80 ± 2.49 49.80 ± 4.58	51.20 ± 2.51 50.30 ± 3.49	69.50 ± 4.57 69.90 ± 5.10
CE VAT PVCC	$86.40 \pm 2.30$ $89.00 \pm 3.87$ $83.40 \pm 4.16$	91.20 ± 1.92 91.20 ± 3.56 82.60 ± 5.22	EMI 85.40 ± 1.95 84.40 ± 2.70 80.00 ± 4.90	96.40 ± 1.82 96.20 ± 1.92 93.20 ± 2.49	CE VAT PVCC	85.80 ± 3.63 85.40 ± 2.88 81.80 ± 5.17	81.80 ± 2.86 83.00 ± 2.55 72.20 ± 6.46	63.80 ± 5.40 60.20 ± 8.98 57.00 ± 9.59	96.20 ± 3.11 97.00 ± 2.83 88.80 ± 6.91	CE VAT PVCC	49.00 ± 4.40 46.80 ± 4.31 52.60 ± 4.02	50.80 ± 2.49 49.80 ± 4.58 48.60 ± 4.68	51.20 ± 2.51 50.30 ± 3.49 46.80 ± 2.02	69.50 ± 4.57 69.90 ± 5.10 66.90 ± 3.49
CE VAT PVCC HGR	86.40 ± 2.30 89.00 ± 3.87 83.40 ± 4.16	91.20 ± 1.92 91.20 ± 3.56 82.60 ± 5.22	EMI 85.40 ± 1.95 84.40 ± 2.70 80.00 ± 4.90	All           96.40 ± 1.82           96.20 ± 1.92           93.20 ± 2.49           95.20 ± 3.27	CE VAT PVCC HGR	85.80 ± 3.63 85.40 ± 2.88 81.80 ± 5.17	81.80 ± 2.86 83.00 ± 2.55 72.20 ± 6.46	63.80 ± 5.40 60.20 ± 8.98 57.00 ± 9.59	All 96.20 ± 3.11 97.00 ± 2.83 88.80 ± 6.91 89.80 ± 4.02	CE VAT PVCC HGR	49.00 ± 4.40 46.80 ± 4.31 52.60 ± 4.02	50.80 ± 2.49 49.80 ± 4.58 48.60 ± 4.68	51.20 ± 2.51 50.30 ± 3.49 46.80 ± 2.02	69.50 ± 4.57 69.90 ± 5.10 66.90 ± 3.49 61.20 ± 5.38
CE VAT PVCC HGR TCGM	$\begin{array}{c} 86.40 \pm 2.30 \\ 89.00 \pm 3.87 \\ 83.40 \pm 4.16 \\ - \\ 94.20 \pm 2.17 \end{array}$	91.20 ± 1.92 91.20 ± 3.56 82.60 ± 5.22 - 93.20 ± 3.27	LMI 85.40 ± 1.95 84.40 ± 2.70 80.00 ± 4.90 - 92.40 ± 1.34	All $96.40 \pm 1.82$ $96.20 \pm 1.92$ $93.20 \pm 2.49$ $95.20 \pm 3.27$ $98.40 \pm 1.14$	CE VAT PVCC HGR TCGM	FAM 85.80 ± 3.63 85.40 ± 2.88 81.80 ± 5.17 - 89.20 ± 2.28	81.80 ± 2.86 83.00 ± 2.55 72.20 ± 6.46 - 87.20 ± 4.32	$ \begin{array}{r}     \text{LMI} \\     63.80 \pm 5.40 \\     60.20 \pm 8.98 \\     57.00 \pm 9.59 \\     - \\     73.60 \pm 4.28 \\ \end{array} $	96.20 ± 3.11           97.00 ± 2.83           88.80 ± 6.91           89.80 ± 4.02           98.20 ± 1.10	CE VAT PVCC HGR TCGM	$49.00 \pm 4.40$ $46.80 \pm 4.31$ $52.60 \pm 4.02$ - 63.60 \pm 4.60	50.80 ± 2.49 49.80 ± 4.58 48.60 ± 4.68 - 61.30 ± 3.40	$51.20 \pm 2.51$ $50.30 \pm 3.49$ $46.80 \pm 2.02$ $-$ $64.70 \pm 1.75$	$69.50 \pm 4.57$ $69.90 \pm 5.10$ $66.90 \pm 3.49$ $61.20 \pm 5.38$ $86.70 \pm 2.97$

Fig. 4. Test accuracies (mean  $\pm$  std. dev. ) on Newsgroups datasets

As shown in Fig. 4, **TCGM** achieves the best classification accuracy for both single classifier and aggregators, especially when the label rate is small. This shows the efficacy of utilizing the cross-modal information during training as compared to others that are unable to utilize the cross-modal information. Moreover, we can achieve further improvement by aggregating classifiers on all modalities, which shows the benefit of aggregating knowledge from different modalities.

#### 6.2 Emotion Recognition

*Dataset* We evaluate our methods on two multi-modal emotion recognition datasets: **IEMOCAP** dataset [7] and **MOSI** dataset [34]. The goal for both datasets is to identify speaker emotions based on the collected videos, audios and texts. The IEMOCAP consists of 151 sessions of recorded dialogues, with 2 speakers per session for a total of 302 videos across the dataset. The MOSI is composed of 93 opinion videos from YouTube movie reviews. We follow the settings in [23] for the data splits of training, validation and test set. For IEMOCAP, we conduct experiments on three different emotions: happy, angry and neutral emotions; for MOSI dataset we consider the binary classification of emotions: positive and negative.

Implementation details We synthesize three label rates for each dataset (the percentage of labeled data points in each modality):  $\{0.5\%, 1\%, 1.5\%\}$  for IEMOCAP and  $\{1\%, 2\%, 3\%\}$  for MOSI. For a fair comparison, we follow architecture setting in [23]. We adopt the modality encoder architectures in [23] as the single modality classifiers for CE and TCGM, while adopting the aggregator on the top of modality encoders for LMF and HGR. Adam with default parameters and learning rate  $\gamma_u = 0.0001, \gamma_l = 0.001$  is used as the optimizer. The batch size is set to 32.

We report the AUC (Area under ROC curve) for the aggregators on all the modalities and single modality classifiers by different methods. We only report the AUC of LMF



Fig. 5. (a,b,c) AUC of Aggregators on happy, angry and neutral emotion recognition on IEMOCAP data: (d, e,f) AUC on text, audio and video modality modality classifiers on IEMOCAP: AUC of other composition of (modality, emotion) are listed in supplementary material. (g,h,i) AUC on MOSI data: AUC of (g) Aggregators on all modalities and single classifiers on (h) Video modality (i) Audio modality.

and HGR on all modalities since they do not have single modality classifiers. For single modality classifiers, we show results on the text modality on happy emotion (d), audio modality on neutral emotion (e) the video modality on angry emotion on IEMOCAP; and (h) the video modality (i) the audio modality on MOSI. Please refer to supplementary material for complete experimental result. As shown in Figure 5, aggregators trained by TCGM outperform all the baselines given only tiny fractions of labeled data. **TCGM** improves the AUC of the single modality classifiers significantly, which shows the efficacy of utilizing the cross-modal information during the training of our method. As label rates continue to grow, the advantage of our method over CE decreases since more information is provided for CE to learn the ground-truth label.

Our method also outperforms other methods in terms of the prediction based on all the modalities, especially when the label rate is small. This shows the superiority of our method when dealing with a limited amount of annotations.

#### 6.3 Disease prediction of Alzheimer's Disease

*Dataset* Early prediction of Alzheimer's Disease (AD) is attracting increasing attention since it is irreversible and very challenging. Besides, due to privacy issues and high collecting costs, an efficient classifier with limited labeled data is desired. To validate the effectiveness of our method on this challenging task, we only keep labels of a limited percentage of data, which is obtained from the most popular **Alzheimer's Disease Neuroimaging Initiative (ADNI)** dataset<sup>8</sup>, with 3D images sMRI and PET. DARTEL VBM pipeline [2] is implemented to pre-process the sMRI data, and then images of PET were reoriented into a standard  $91 \times 109 \times 91$  voxel image grid in MNI152 space, which is same with sMRIs'. To limit the size of images, only the hippocampus on both sides are extracted as input in the experiments. We denote subjects that convert to Alzheimer's disease (MCI<sub>c</sub>) as AD, and subjects that remain stable (MCI<sub>s</sub>) as NC (Normal Control). Our dataset contains 300 samples in total, with 144 AD and 156 NC. We randomly choose 60% for training, 20% for validation and 20% for testing stage.

*Implementation details* We synthesize two different label rates (the percentage of labeled data points): {10%, 50%}. DenseNet is used as the classifier. Two 3D convolutional layers with the kernel size  $3 \times 3 \times 3$  are adopted to replace the first 2D convolutional layers with the kernel size  $7 \times 7$ . We use four dense blocks with the size of (6, 12, 24, 16). To preserve more low-level local information, we discard the first max-pooling layer that follows after the first convolution layer. Adam with default parameters and learning rate  $\gamma_u = \gamma_l = 0.001$  are used as the optimizer during training. We set Batch Size as only 12 due to the large memory usage of 3D images. Random crop of  $64 \times 64 \times 64$ , random flip and random transpose are applied as data augmentation.

	sMRI	PET	All		sMRI	PET	All		
CE	$54.48 \pm 8.76$	$58.62 \pm 3.23$	$62.07 \pm 2.72$	CE	$63.79 \pm 3.66$	$60.69 \pm 2.83$	$65.52 \pm 3.66$		
HGR	-	-	$60.34 \pm 5.59$	HGR	-	_	$66.22 \pm 4.16$		
TCGM	$63.44 \pm 2.83$	$62.41 \pm 3.53$	$65.17 \pm 1.44$	TCGM	$65.86 \pm 3.09$	$62.76 \pm 1.97$	$67.24 \pm 2.10$		
	ADNI, Lat	el Rate 10%		ADNI, Label Rate 50%					

Fig. 6. Test accuracies (mean  $\pm$  std.) on ADNI dataset

Figure 6 shows the accuracy of classifiers for each modality and the aggregator. Our method **TCGM** outperforms the baseline methods in all settings especially when the label rate is small, which is desired since it is costly to label data.

To further illustrate the advantage of our model over others in terms of leveraging the knowledge of another modality, we visualize two MCI<sub>c</sub>s, denoted as MCI<sub>c</sub><sup>1</sup> and MCI<sub>c</sub><sup>2</sup>, which are mistakenly classified as NC by **CE**'s classifier for sMRI and PET modality, respectively. The volume and standardized uptake value (SUV) (a measurement of the degree of metabolism), whose information are respectively contained by sMRI and PET data, are linearly mapped to the darkness of the red and blue. Darker color implies smaller volume and SUV, i.e., more probability of being AD. As shown in Figure 7,

<sup>&</sup>lt;sup>8</sup> www.loni.ucla.edu/ADNI



**Fig. 7.** Volume (sMRI, top line) and SUV (PET, bottom line) of  $MCI_c^1$ ,  $MCI_c^2$ , AD and NC. Darker color implies smaller volume and SUV, i.e., more probability of being AD.

the volume (SUV) of  $MCI_c^1$  ( $MCI_c^2$ ) is similar to NC, hence it is reasonable for **CE** to mistakenly classify it by only using the information of volume (SUV). In contrast, **TCGM** for each modality can correctly classify both cases as AD, which shows the better learning of the information intersection (i.e., the ground truth) during training, facilitated by the leverage of knowledge from another modality.

### 7 Conclusion

In this paper, we propose an information-theoretic framework on multi-modal data, Total Correlation Gain Maximization (TCGM), in the scenario of semi-supervised learning. Specifically, we learn to infer the ground truth labels shared by all modalities by maximizing the total correlation gain. Conditioning on a common assumption that all modalities are independent given the ground truth label, it can be theoretically proved our method can learn the Bayesian posterior classifier for each modality and the Bayesian posterior aggregator for all modalities. Extensive experiments on Newsgroup, IEMOCAP, MOSI and ADNI datasets are conducted and achieve promising results, which demonstrates the benefit and utility of our framework.

# Acknowledgement

Yizhou Wang's work is supported by MOST-2018AAA0102004, NSFC-61625201, DFG TRR169/NSFC Major International Collaboration Project "Crossmodal Learning". Yuqing Kong's work is supported by Science and Technology Innovation 2030 – "The New Generation of Artificial Intelligence" Major Project No.2018AAA0100901, China. Thanks to Xinwei Sun's girlfriend Yue Cao, for all her love and support.

# References

- Achille, A., Soatto, S.: Emergence of invariance and disentanglement in deep representations. The Journal of Machine Learning Research 19(1), 1947–1980 (2018) 6
- Ashburner, J.: A fast diffeomorphic image registration algorithm. Neuroimage 38(1), 95–113 (2007) 10, 13
- 3. Balcan, M.F., Blum, A., Yang, K.: Co-training and expansion: Towards bridging theory and practice. In: Advances in neural information processing systems. pp. 89–96 (2005) 3
- Belghazi, M.I., Baratin, A., Rajeswar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, R.D.: Mine: mutual information neural estimation. arXiv preprint arXiv:1801.04062 (2018) 4, 5
- Bisson, G., Grimal, C.: Co-clustering of multi-view datasets: a parallelizable approach. In: 2012 IEEE 12th International Conference on Data Mining. pp. 828–833. IEEE (2012) 10
- Blum, A., Mitchell, T.: Combining labeled and unlabeled data with co-training. In: Proceedings of the eleventh annual conference on Computational learning theory. pp. 92–100. ACM (1998) 2, 3, 6
- Busso, C., Bulut, M., Lee, C.C., Kazemzadeh, A., Provost, E.M., Kim, S., Chang, J.N., Lee, S., Narayanan, S.S.: Iemocap: interactive emotional dyadic motion capture database. Language Resources and Evaluation 42, 335–359 (2008) 11
- Cao, P., Xu, Y., Kong, Y., Wang, Y.: Max-mig: an information theoretic approach for joint learning from crowds. ICLR2019 (2018) 4, 7
- Chandar, S., Khapra, M.M., Larochelle, H., Ravindran, B.: Correlational neural networks. Neural computation 28(2), 257–285 (2016) 3
- Chang, X., Xiang, T., Hospedales, T.M.: Scalable and effective deep cca via soft decorrelation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1488–1497 (2018) 3
- Cheng, Y., Zhao, X., Huang, K., Tan, T.: Semi-supervised learning and feature evaluation for rgb-d object recognition. Computer Vision and Image Understanding 139, 149–160 (2015) 3
- Christoudias, C.M., Saenko, K., Morency, L.P., Darrell, T.: Co-adaptation of audio-visual speech and gesture classifiers. In: Proceedings of the 8th international conference on Multimodal interfaces. pp. 84–91. ACM (2006) 3
- Dasgupta, S., Littman, M.L., McAllester, D.A.: Pac generalization bounds for co-training. In: Advances in neural information processing systems. pp. 375–382 (2002) 2, 6
- Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. arXiv preprint arXiv:1808.06670 (2018) 4
- Huang, S.L., Suh, C., Zheng, L.: Euclidean information theory of networks. IEEE Transactions on Information Theory 61(12), 6795–6814 (2015) 5
- Hussain, S.F., Grimal, C., Bisson, G.: An improved co-similarity measure for document clustering. In: International Conference on Machine Learning and Applications (December 2010) 10
- Jones, R.: Learning to extract entities from labeled and unlabeled text. Ph.D. thesis, Citeseer (2005) 3
- Kim, D.H., Lee, M.K., Choi, D.Y., Song, B.C.: Multi-modal emotion recognition using semisupervised learning and multiple neural networks in the wild. In: Proceedings of the 19th ACM International Conference on Multimodal Interaction. pp. 529–535. ACM (2017) 3
- Kong, Y., Schoenebeck, G.: Water from two rocks: Maximizing the mutual information. In: Proceedings of the 2018 ACM Conference on Economics and Computation. pp. 177–194. ACM (2018) 4, 7
- Leskes, B.: The value of agreement, a new boosting algorithm. In: International Conference on Computational Learning Theory. pp. 95–110. Springer (2005) 2, 6
- Levin, A., Viola, P.A., Freund, Y.: Unsupervised improvement of visual detectors using co-training. In: ICCV. vol. 1, p. 2 (2003) 3
- Lewis, D.D.: Naive (bayes) at forty: The independence assumption in information retrieval. In: European conference on machine learning. pp. 4–15. Springer (1998) 2, 6

- 16 Sun et al.
- 23. Liu, Z., Shen, Y., Lakshminarasimhan, V.B., Liang, P.P., Zadeh, A., Morency, L.P.: Efficient low-rank multimodal fusion with modality-specific factors. In: ACL (2018) 10, 11
- Ng, A.Y., Jordan, M.I., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: NIPS (2001) 10
- Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., Ng, A.Y.: Multimodal deep learning. In: Proceedings of the 28th international conference on machine learning (ICML-11). pp. 689–696 (2011) 2, 3
- Nguyen, X., Wainwright, M.J., Jordan, M.I.: Estimating divergence functionals and the likelihood ratio by convex risk minimization. IEEE Transactions on Information Theory 56(11), 5847–5861 (2010) 5
- Sohn, K., Shang, W., Lee, H.: Improved multimodal deep learning with variation of information. In: Advances in Neural Information Processing Systems. pp. 2141–2149 (2014) 2, 3
- 28. Studenỳ, M., Vejnarová, J.: The multiinformation function as a tool for measuring stochastic dependence. In: Learning in graphical models, pp. 261–297. Springer (1998) 2, 4
- 29. Takeru Miyato, Shin-ichi Maeda, M.K.S.I.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. In: IEEE transactions on pattern analysis and machine intelligence (2018). 11
- Veličković, P., Fedus, W., Hamilton, W.L., Liò, P., Bengio, Y., Hjelm, R.D.: Deep graph infomax. arXiv preprint arXiv:1809.10341 (2018) 4
- Wang, L., Wu, J., Huang, S.L., Zheng, L., Xu, X., Zhang, L., Huang, J.: An efficient approach to informative feature extraction from multimodal data. arXiv preprint arXiv:1811.08979 (2018) 3, 10
- Xu, Y., Zhao, S., Song, J., Stewart, R.J., Ermon, S.: A theory of usable information under computational constraints. ArXiv abs/2002.10689 (2020) 4
- Yang, Y., Zhan, D.C., Sheng, X.R., Jiang, Y.: Semi-supervised multi-modal learning with incomplete modalities. In: IJCAI. pp. 2998–3004 (2018) 10, 11
- Zadeh, A., Zellers, R., Pincus, E., Morency, L.P.: Mosi: Multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos. ArXiv abs/1606.06259 (2016) 11