

# Recognizable Information Bottleneck

Yilin Lyu<sup>1</sup>, Xin Liu<sup>1</sup>, Mingyang Song<sup>1</sup>, Xinyue Wang<sup>1</sup>,  
Yaxin Peng<sup>2</sup>, Tiejong Zeng<sup>3</sup> and Liping Jing<sup>1\*</sup>

<sup>1</sup>Beijing Key Lab of Traffic Data Analysis and Mining, Beijing Jiaotong University

<sup>2</sup>Department of Mathematics, School of Science, Shanghai University

<sup>3</sup>The Chinese University of Hong Kong

{yilinlyu, xin.liu, mingyang.song, xinyuewang, lpjing}@bjtu.edu.cn,  
yaxin.peng@shu.edu.cn, zeng@math.cuhk.edu.hk

## Abstract

Information Bottlenecks (IBs) learn representations that generalize to unseen data by information compression. However, existing IBs are practically unable to guarantee generalization in real-world scenarios due to the vacuous generalization bound. The recent PAC-Bayes IB uses information complexity instead of information compression to establish a connection with the mutual information generalization bound. However, it requires the computation of expensive second-order curvature, which hinders its practical application. In this paper, we establish the connection between the recognizability of representations and the recent functional conditional mutual information ( $f$ -CMI) generalization bound, which is significantly easier to estimate. On this basis we propose a Recognizable Information Bottleneck (RIB) which regularizes the recognizability of representations through a recognizability critic optimized by density ratio matching under the Bregman divergence. Extensive experiments on several commonly used datasets demonstrate the effectiveness of the proposed method in regularizing the model and estimating the generalization gap.

## 1 Introduction

Learning representations that are generalized to unseen data is one of the most fundamental goals for machine learning. The IB method [Tishby *et al.*, 2000] is one of the prevailing paradigms. It centers on learning representations that not only sufficiently retain information about the label, but also compress information about the input. The rationale behind the IB is that the information compression of representations penalizes complex hypotheses and guarantees generalization by a probability generalization bound [Shamir *et al.*, 2010; Vera *et al.*, 2018]. It is expected that one can obtain a lower generalization gap through optimizing the IB training objective. However, in finite and continuous data scenarios, such

as an image classification task with deep neural networks (DNNs), this upper bound is dominated by a prohibitively large constant with any reasonable probability and becomes vacuous, thus making the information compression term negligible [Rodriguez Galvez, 2019]. This deprives IB method of the theoretical guarantee in real-world scenarios.

Recently, motivated by the mutual information (MI)-based generalization bound in the information-theoretic framework [Xu and Raginsky, 2017], Wang *et al.* [2022] proposed to use the *information stored in weights*, i.e., the MI between the input dataset and the output weights of the algorithm, to replace the compression term in the IB objective. It is worth noting, however, that the MI-based generalization bound may still be vacuous due to the fact that the MI can easily go to infinity, e.g., in cases where the learning algorithm is a deterministic function of the input data [Bu *et al.*, 2020; Steinke and Zakyntinou, 2020]. More importantly, it is notoriously difficult to estimate the MI between two high-dimensional random variables, e.g., weights of model and a training dataset. Although one can recover non-vacuous bound by restricting the prior and posterior of the weights to Gaussian distributions [Dziugaite and Roy, 2017; Wang *et al.*, 2022], the computational burdens (from e.g., Monte-Carlo sampling or calculating the second-order curvature information) are still cumbersome.

To address the unbounded nature of MI-based bound, Steinke and Zakyntinou [2020] proposed the celebrated conditional mutual information (CMI) generalization bound, which normalizes each datum to one bit, thus is always bounded. The intuition is that measuring information that “recognizes” an input sample is much simpler than measuring information that “reconstructs” an input sample. Harutyunyan *et al.* [2021] further extended this idea to a black-box algorithm setting and proposed a functional CMI ( $f$ -CMI) generalization bound, which significantly reduces the computational cost by computing the CMI with respect to the low-dimensional predictions rather than the high-dimensional weights. This inspires us to consider the DNN encoder as a black box and the representation as the output predictions, which coincides with the original IB setting. Unfortunately, the computation of the  $f$ -CMI bound requires multiple runs on multiple draws of data, there is not yet a practical optimization algorithm to incorporate it into the training of the models.

\*Corresponding author.

Appendices are available at: <https://arxiv.org/pdf/2304.14618>.

In this paper, we first formulate the *recognizability* from the perspective of *binary hypothesis testing* (BHT) and show the connection between the *recognizability of representations* and the *f*-CMI generalization bound. Based on the reasoning that lower recognizability leads to better generalization, we propose a new IB objective, called the Recognizable Information Bottleneck (RIB). A recognizability critic is exploited to optimize the RIB objective using *density ratio matching* under the Bregman divergence. Extensive experiments on several commonly used datasets demonstrate the effectiveness of the RIB in controlling generalization and the potential of estimating the generalization gap for the learned models using the recognizability of representations.

Our main contributions can be concluded as follows:

- We define a formal characterization of recognizability and then show the connection between the recognizability of representations and the *f*-CMI bound.
- We propose the RIB objective to regularize the recognizability of representations.
- We propose an efficient optimization algorithm of RIB which exploits a recognizability critic optimized by density ratio matching.
- We empirically demonstrate the regularization effects of RIB and the ability of estimating the generalization gap by the recognizability of representations.

### 1.1 Other Related Work

**Information bottlenecks.** There have been numerous variants of IB [Alemi *et al.*, 2017; Kolchinsky *et al.*, 2019; Gálvez *et al.*, 2020; Pan *et al.*, 2021; Yu *et al.*, 2021] and have been utilized on a wide range of machine learning tasks [Du *et al.*, 2020; Lai *et al.*, 2021; Li *et al.*, 2022]. However, one pitfall of IBs is that there may be no causal connection between information compression and generalization in real-world scenarios [Saxe *et al.*, 2018; Rodriguez Galvez, 2019; Dubois *et al.*, 2020]. Wang *et al.* [2022] addressed this issue by introducing an information complexity term to replace the compression term, which turns out to be the Gibbs algorithm stated in the PAC-Bayesian theory [Catoni, 2007; Xu and Raginsky, 2017].

**Information-theoretical generalization bounds.** The MI between the input and output of a learning algorithm has been widely used to bound the generalization error [Xu and Raginsky, 2017; Pensia *et al.*, 2018; Negrea *et al.*, 2019; Bu *et al.*, 2020]. To tackle the intrinsic defects of MI, Steinke and Zakyntinou [2020] proposed the CMI generalization bound, then Haghifam *et al.* [2020] proved that the CMI-based bounds are tighter than MI-based bounds. Zhou *et al.* [2022] made the CMI bound conditioned on an individual sample and obtained a tighter bound. The above bounds are based on the output of the algorithm (e.g. weights of the model) rather than information of predictions which is much easier to estimate in practice [Harutyunyan *et al.*, 2021].

## 2 Preliminaries

### 2.1 Notation and Definitions

We use capital letters for random variables, lower-case letters for realizations and calligraphic letters for domain. Let  $X$  and  $Y$  be random variables defined on a common probability space, the mutual information between  $X$  and  $Y$  is defined as  $I(X; Y) = D_{\text{KL}}(P_{X,Y} \| P_X \otimes P_Y)$ , where  $D_{\text{KL}}$  is the Kullback-Leibler (KL) divergence. The conditional mutual information is defined as  $I(X; Y|Z) = D_{\text{KL}}(P_{X,Y|Z} \| P_{X|Z} P_{Y|Z} | P_Z) = \mathbb{E}_{z \sim P_Z} [I(X; Y|Z = z)]$ . A random variable  $X$  is  $\sigma$ -subgaussian if  $\mathbb{E}[\exp(\lambda(X - \mathbb{E}(X)))] \leq \exp(\lambda^2 \sigma^2 / 2)$  for all  $\lambda \in \mathbb{R}$ .

Throughout this paper, we consider the supervised learning setting. There is an instance random variable  $X$ , a label random variable  $Y$ , an unknown data distribution  $\mathcal{D}$  over the  $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$  and a training dataset  $S = (Z_1, Z_2, \dots, Z_n) \sim \mathcal{D}^n$  consisting of  $n$  i.i.d. samples drawn from the data distribution. There is a learning algorithm  $A : \mathcal{Z}^n \rightarrow \mathcal{W}$  and an encoder function  $f_\theta : \mathcal{X} \rightarrow \mathcal{T}$  from its class  $\mathcal{F} := \{T = \mathbb{E}_X[\mathbf{1}[t = f_\theta(X)]] : \theta \in \mathcal{W}\}$ . We assume that  $\mathcal{T}$  is of finite size and consider  $T$  as the representation. We define two expressions of generalization, one from the “weight” perspective and one from the “representation” perspective:

1. Given a loss function  $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ , the true risk of the algorithm  $A$  w.r.t.  $\ell$  is  $\mathcal{L}(w, S) = \mathbb{E}_{z' \sim \mathcal{D}} \ell(w, z')$  and the empirical risk of the algorithm  $A$  w.r.t.  $\ell$  is  $\mathcal{L}_{\text{emp}}(w, S) = \frac{1}{n} \sum_{i=1}^n \ell(w, z_i)$ . The generalization error is defined as  $\text{gen}(w, S) = \mathcal{L}(w, S) - \mathcal{L}_{\text{emp}}(w, S)$ ;
2. Given a loss function  $\ell : \mathcal{T} \times \mathcal{Y} \rightarrow \mathbb{R}$ , the true risk of the algorithm  $A$  w.r.t.  $\ell$  is  $\mathcal{L}(T, Y) = \mathbb{E}_{z' \sim \mathcal{D}} \ell(t', y')$  and the empirical risk of the algorithm  $A$  w.r.t.  $\ell$  is  $\mathcal{L}_{\text{emp}}(T, Y) = \frac{1}{n} \sum_{i=1}^n \ell(t_i, y_i)$ . The generalization error is defined as  $\text{gen}(T, Y) = \mathcal{L}(T, Y) - \mathcal{L}_{\text{emp}}(T, Y)$ .

### 2.2 Generalization Guarantees for Information Bottlenecks

The IB method [Tishby *et al.*, 2000] wants to find the optimal representation by minimizing the following Lagrangian:

$$\mathcal{L}_{\text{IB}} = \mathcal{L}_{\text{emp}}(T, Y) + \beta I(T; X), \tag{1}$$

where  $\beta$  is the Lagrange multiplier. The  $I(T; X)$  in Eq. (1) can be seen as a compression term that regularizes the complexity of  $T$ . It has been demonstrated to upper-bound the generalization error [Shamir *et al.*, 2010; Vera *et al.*, 2018]. Concretely, for any given  $\delta > 0$ , with probability at least  $1 - \delta$  it holds that

$$|\text{gen}(T, Y)| \leq \mathcal{O}\left(\frac{\log n}{\sqrt{n}}\right) \sqrt{I(T; X)} + C_\delta,$$

where  $C_\delta = \mathcal{O}(|\mathcal{T}|/\sqrt{n})$  is a constant that depends on the cardinality of the space of  $T$ . The bound implies that **the less information the representations can provide about the instances, the better the generalization of the algorithm**. However, since  $|\mathcal{T}|$  will become prohibitively large with any reasonable choice of  $\delta$ , this bound is actually vacuous in practical scenarios [Rodriguez Galvez, 2019].

A recent approach PAC-Bayes IB (PIB) [Wang *et al.*, 2022] tried to amend the generalization guarantee of IB by replacing the compression term with the information stored in weights, which leads to the following objective:

$$\mathcal{L}_{\text{PIB}} = \mathcal{L}_{\text{emp}}(w, S) + \beta I(w; S).$$

The rationale behind the PIB is stem from the input-output mutual information generalization bound which can also be derived from the PAC-Bayesian perspective [Xu and Raginsky, 2017; Banerjee and Montúfar, 2021]. Concretely, assume that the loss function  $\ell(w, Z)$  is  $\sigma$ -subgaussian for all  $w \in \mathcal{W}$ , the expected generalization error can be upper-bounded by

$$|\mathbb{E}[\text{gen}(w, S)]| \leq \sqrt{\frac{2\sigma^2}{n} I(W; S)}.$$

Intuitively, the bound implies that **the less information the weight can provide about the input dataset, the better the generalization of the algorithm**. However, as we mentioned above, the MI-based bound could be vacuous without proper assumptions and is hard to estimate for two high-dimensional variables such as the weight  $W$  and the training dataset  $S$ .

### 2.3 The CMI-based Generalization Bounds

The CMI-based bounds [Steinke and Zakynthinou, 2020; Harutyunyan *et al.*, 2021] characterize the generalization by measuring the ability of recognizing the input data given the algorithm output. This can be accomplished by recognizing the training samples from the “ghost” (possibly the test) samples. Formally, consider a supersample  $\tilde{Z} \in \mathcal{Z}^{n \times 2}$  constructed from  $n$  input samples mixed with  $n$  ghost samples. A selector variable  $U \sim \text{Unif}(\{0, 1\}^n)$  selects the input samples from the supersample  $\tilde{Z}$  to form the training set  $S = \tilde{Z}_U = (\tilde{Z}_{i, U_i+1})_{i=1}^n$ . Given a loss function  $\ell(w, z) \in [0, 1]$  for all  $w \in \mathcal{W}$  and all  $z \in \mathcal{Z}$ , Steinke and Zakynthinou [2020] bound the expected generalization error by the CMI between the weight  $W$  and the selector variable  $U$  given the supersample  $\tilde{Z}$ :

$$|\mathbb{E}[\text{gen}(w, S)]| \leq \sqrt{\frac{2}{n} I(W; U | \tilde{Z})}. \quad (2)$$

The intuition is that **the less information the weight can provide to recognize input samples from their ghosts, the better the generalization of the algorithm**. In this way, the algorithm that reveals a datum has only CMI of one bit, bounding the CMI by  $n \log 2$ , while the MI may be infinite. Harutyunyan *et al.* [2021] extended this idea to a black-box algorithm setting which measures the CMI with respect to the predictions of the model rather than the weight. It is able to further extend the prediction domain to the representation domain, which yields a setting similar to the original IB. Given a loss function  $\ell(t, y) \in [0, 1]$  for all  $t \in \mathcal{T}$  and all  $z \in \mathcal{Z}$ , the expected generalization error is upper-bounded by the following  $f$ -CMI bound:

$$|\mathbb{E}[\text{gen}(T, Y)]| \leq \sqrt{\frac{2}{n} I(T; U | \tilde{Z})}. \quad (3)$$

From the Markov chain  $U \rightarrow W \rightarrow T$  and the data processing inequality, it follows that  $I(T; U | \tilde{Z}) \leq I(W; U | \tilde{Z})$ ,




Figure 1: Illustrations of recognizability.

which indicates the bound in Eq. (3) is tighter than the bound in Eq. (2). Similarly, we can learn from the bound that **the less information the representations can provide to recognize input samples from their ghosts, the better the generalization of the algorithm**. Since the representation  $T$  and the selector variable  $U$  are relatively lower dimensional than the weight  $W$  and the dataset  $S$ , respectively, estimating the  $f$ -CMI bound is far more efficient than the previous bounds.

## 3 Methodology

Although the  $f$ -CMI bound is theoretically superior to the generalization bounds of existing IBs, it has so far only been used as a risk certificate for the learned model [Harutyunyan *et al.*, 2021]. Since computing the  $f$ -CMI bound requires multiple runs on multiple draws of data, there is no immediate way to incorporate it into the training of deep models. In this section, we first give a formal characterization of the recognizability and describe how to use the recognizability of representations to approximate the  $f$ -CMI bound. This motivates us to propose the RIB objective and an efficient optimization algorithm using a recognizability critic.

### 3.1 Recognizability

In order to characterize the recognizability, we first review some basic concepts of binary hypothesis testing (BHT) [Levy, 2008]. Assume that there are two possible hypotheses,  $H_0$  and  $H_1$ , corresponding to two possible distributions  $P$  and  $Q$  on a space  $\mathcal{X}$ , respectively,

$$H_0 : X \sim P \quad H_1 : X \sim Q.$$

A binary hypothesis test between two distributions is to choose either  $H_0$  or  $H_1$  based on the observation of  $X$ . Given a decision rule  $\Xi : \mathcal{X} \rightarrow \{0, 1\}$  such that  $\Xi = 0$  denotes the test chooses  $P$ , and  $\Xi = 1$  denotes the test chooses  $Q$ . Then the recognizability can be defined as follows:

**Definition 1** (Recognizability). *Let  $H_0, H_1, P, Q, X, \Xi$  be defined as above. The recognizability between  $P$  and  $Q$  w.r.t.  $\Xi$  for all randomized tests  $P_{\Xi|X}$  is defined as*

$$\mathfrak{R}_{\Xi}(P, Q) := \left| \left\{ (\Pr(\Xi = 1|H_0), \Pr(\Xi = 1|H_1)) \mid \forall \Xi \sim P_{\Xi|X} \right\} \right|,$$

where  $\Pr(\Xi = 1|H_0)$  is the probability of error given  $H_1$  is true and  $\Pr(\Xi = 1|H_1)$  is the probability of success given  $H_0$  is true.

**Remark.** In words, the recognizability is the area of the achievable region for all tests drawn from  $P_{\Xi|X}$  (see Fig. 1a). For BHT, one would like to select a test such that  $\Pr(\Xi = 1|H_0)$  is as close to zero as possible while  $\Pr(\Xi = 1|H_1)$  is as close to one as possible. However, the extreme case, i.e.,  $(\Pr(\Xi = 1|H_0) = 0, \Pr(\Xi = 1|H_1) = 1)$  or  $(0, 1)$ , is achievable only if the hypotheses  $H_0$  and  $H_1$  are completely mutually exclusive, that is,  $P$  is mutually singular w.r.t.  $Q$ , which indicates that the recognizability  $\mathfrak{R}_{\Xi}(P, Q) = 1$  (Fig. 1b). Similarly, when  $H_0$  and  $H_1$  are equivalent, that is,  $P$  and  $Q$  are identical, we have the recognizability  $\mathfrak{R}_{\Xi}(P, Q) = 0$  (Fig. 1c).

By the celebrated Neyman-Pearson lemma, we have that log-likelihood ratio testing (LRT) is the optimal decision rule, which means all the points in the achievable region are attained by LRT [Neyman *et al.*, 1933]. It is also well-known that the upper boundary between the achievable and unachievable region is the receiver operating characteristic (ROC) of LRT [Levy, 2008]. Therefore, given a LRT decision rule  $\xi : \mathcal{X} \rightarrow [0, 1]$ , the recognizability w.r.t.  $\xi$  can be computed by

$$\mathfrak{R}_{\xi}(P, Q) = 2 \times \text{AUCROC}_{\xi} - 1, \quad (4)$$

where  $\text{AUCROC}_{\xi}$  is the area under the ROC curve of  $\xi$ .

The formulation of recognizability offers us some fresh insights to consider generalization from the perspective of BHT, which leads to a new approach to regularization. Concretely, we rewrite the  $f$ -CMI term in Eq. (3) as

$$\begin{aligned} I(T; U|\tilde{Z}) &= \mathbb{E}_{\tilde{z} \sim \tilde{Z}} D_{\text{KL}}(P_{T,U|\tilde{Z}=\tilde{z}} \| P_{T|\tilde{Z}=\tilde{z}} \otimes P_U) \\ &= \mathbb{E}_{\tilde{z} \sim \tilde{Z}} D_{\text{KL}}(P_{T|U, \tilde{Z}=\tilde{z}} \| P_{T|\tilde{Z}=\tilde{z}}), \end{aligned} \quad (5)$$

where Eq. (5) follows by the uniformity of  $U$ . Consider a realization of supersample  $\tilde{z}$ . Let  $t_1, t_2, \dots, t_n$  be a sequence drawn i.i.d. according to  $P_{T|U, \tilde{Z}=\tilde{z}}$ , then by the weak law of large numbers we have that

$$\frac{1}{n} \sum_{i=1}^n \log \frac{P_{T|U, \tilde{Z}=\tilde{z}}(t_i)}{P_{T|\tilde{Z}=\tilde{z}}(t_i)} \rightarrow D_{\text{KL}}(P_{T|U, \tilde{Z}=\tilde{z}} \| P_{T|\tilde{Z}=\tilde{z}}).$$

This suggests that the log-likelihood ratio between the distributions  $P_{T|U, \tilde{Z}=\tilde{z}}$  and  $P_{T|\tilde{Z}=\tilde{z}}$  amounts to the KL divergence, and eventually an estimate of  $I(T; U|\tilde{z})$ , which is the single-draw of the  $f$ -CMI bound. Thus, the more similar these two distributions are, the more difficult it is for us to determine which hypothesis  $t_i$  originates from by LRT, namely, the lower the recognizability; we can also verify, in turn, that the KL divergence is minimal when we achieve the lowest recognizability (Fig. 1c). To make this explicit, we theoretically establish the connection between recognizability of representations and the  $f$ -CMI bound through the following theorem.

**Theorem 1.** *Let  $\tilde{z}$  be a realization of supersample  $\tilde{Z}$ ,  $T$  be the representation variable defined as above,  $\xi : \mathcal{T} \rightarrow [0, 1]$  be the LRT decision rule and  $U \sim \text{Unif}(\{0, 1\}^n)$ . The recognizability of representations  $\mathfrak{R}_{\xi}(P_{T|\tilde{Z}=\tilde{z}}, P_{T|U, \tilde{Z}=\tilde{z}})$  is upper bound by  $I(T; U|\tilde{z}) + \log \frac{e}{2}$ .*

*Proof.* See Appendix A.  $\square$

Inspired by the above insights, we propose the RIB to constrain the model to learn representations having low recognizability, which utilizes a log-likelihood ratio estimator that we name *recognizability critic*. We will detail it in the next subsection.

### 3.2 Recognizable Information Bottleneck

We first define a new objective function that takes the recognizability of representations into consideration. Let  $\tilde{z}$  be a realization of  $\tilde{Z}$  which is composed of a training set  $S$  of size  $n$  and a ghost set  $\tilde{S}$  of the same size,  $U \sim \text{Unif}(\{0, 1\}^n)$  be a random variable indicating the set from which the instance of representation comes. Note that in general the ghost set is built from the validation or test (if available) set with dummy labels, and we will experiment later with ghost from different sources. Now we can write a new IB objective called the Recognizable Information Bottleneck (RIB) as follows:

$$\mathcal{L}_{\text{RIB}} = \mathcal{L}_{\text{emp}}(T, Y) + \beta I(T; U|\tilde{z}), \quad (6)$$

where the first term enforces empirical risk minimization, the second term regularizes the recognizability of representations and  $\beta$  is the trade-off parameter.

To optimize the second term of Eq. (6), we first train a log-likelihood ratio estimator  $V_{\phi} : \mathcal{T}^2 \rightarrow \mathbb{R}$  that we call the recognizability critic. It takes a pair of representations as input, one obtained from the training set  $S$  samples and the other from the ghost set  $\tilde{S}$ . The pair of representations is then randomly concatenated and treated as a single draw from the joint distribution  $P_{T,U|\tilde{Z}=\tilde{z}}$ , since the probability of occurrence of the possible outcomes is consistent. Next we train the recognizability critic by maximizing the lower bound on the Jensen-Shannon divergence as suggested in [Hjelm *et al.*, 2019; Poole *et al.*, 2019] by the following objective:

$$\begin{aligned} \max_{\phi} \mathbb{E}_{\tilde{t} \sim P_{T,U|\tilde{Z}=\tilde{z}}} [-\text{sp}(-V_{\phi}(\tilde{t}))] \\ - \mathbb{E}_{\tilde{t} \sim P_{T,U|\tilde{Z}=\tilde{z}}} [\text{sp}(V_{\phi}(\tilde{t}))], \end{aligned} \quad (7)$$

where  $\text{sp}(x) = \log(1 + e^x)$  is the softplus function. Once the critic  $V$  reaches the maximum, we can obtain the density ratio by

$$R(\tilde{t}) = \frac{P_{T|U, \tilde{Z}=\tilde{z}}(\tilde{t})}{P_{T|\tilde{Z}=\tilde{z}}(\tilde{t})} \approx \exp(V_{\phi}(\tilde{t})).$$

Since the lowest recognizability achieves when the two distributions are identical, that is, the density ratio  $R^* \equiv 1$ , to regularize the representations, we use a surrogate loss motivated by *density-ratio matching* [Sugiyama *et al.*, 2012], which optimizes the Bregman divergence between the optimal density ratio  $R^*$  and the estimated density ratio  $R(\tilde{t})$  defined as

$$D_{\text{BR}_F}(R^* \| R(\tilde{t})) := F(R^*) - F(R(\tilde{t})) - \partial F(R(\tilde{t}))(R^* - R(\tilde{t})). \quad (8)$$

Here,  $F$  is a differentiable and strictly convex function and  $\partial F$  is its derivative. We can derive different measures by using different  $F$ . In the experiments we use  $F(R) = R \log R - (1 + R) \log(1 + R)$ , at which point the Bregman divergence is reduced to the binary KL divergence. More examples and comparison results are given in Appendix C.

---

**Algorithm 1** Optimization of Recognizable Information Bottleneck (RIB)

---

**Input:** training set  $S = \{(x_i, y_i)\}_{i=1}^N$ , ghost set  $\bar{S} = \{\bar{x}_i\}_{i=1}^N$ , encoder  $f_\theta$ , initial encoder parameter  $\theta_0$ , recognizability critic  $V_\phi$ , initial recognizability critic parameter  $\phi_0$ , batch size  $B$ , learning rate  $\eta$

**Output:** trained encoder parameter  $\theta$ , trained recognizability critic parameter  $\phi$

- 1: initialize:  $\theta \leftarrow \theta_0, \phi \leftarrow \phi_0$
  - 2: **while** not converged **do**
  - 3:   sample training mini-batch  $\{(x_i, y_i)\}_{i=1}^B \sim S$
  - 4:   sample ghost mini-batch  $\{\bar{x}_i\}_{i=1}^B \sim \bar{S}$
  - 5:   sample  $\{u_i\}_{i=1}^B \sim \text{Unif}(\{0, 1\})$
  - 6:   compute  $\{t_i = f_\theta(x_i)\}_{i=1}^B, \{\bar{t}_i = f_\theta(\bar{x}_i)\}_{i=1}^B$
  - 7:   set  $\tilde{t} = \{\tilde{t}_i\}_{i=1}^B$ , where
 
$$\tilde{t}_i = \begin{cases} t_i \circ \bar{t}_i, & u_i = 0 \\ \bar{t}_i \circ t_i, & u_i = 1 \end{cases}$$
  - 8:   compute gradient  $g_\phi$  by Eq. (7) and update  $\phi$ :  
 $\phi \leftarrow \phi + \eta g_\phi$
  - 9:   compute gradient  $g_\theta$  by Eq. (9) and update  $\theta$ :  
 $\theta \leftarrow \theta - \eta g_\theta$
  - 10: **end while**
  - 11: **return**  $(\theta, \phi)$
- 

Now the objective function for the encoder  $f_\theta$  can be written as

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(t_i, y_i) + \beta D_{\text{BR}}(R^* \| R(\tilde{t}_i)). \quad (9)$$

The complete optimization algorithm is described in Algorithm 1. Note that although it is also feasible to directly minimize Eq. (7) with respect to the encoder parameter  $\theta$  which leads to a minimax game similar to GAN [Goodfellow *et al.*, 2014; Nowozin *et al.*, 2016], we empirically find that this may cause training instability especially when data is scarce (as detailed in the next section). We consider this is due to the randomness of  $U$  which leads to large variance gradients and hinders the convergence of the model, while we utilize a fixed ratio for optimization thus it is not affected.

When the recognizability critic finishes training, we can obtain the class probability from the estimated density ratio via the sigmoid function, which gives us the following decision rules:

$$\hat{\xi}(\tilde{t}) = \frac{1}{1 + \exp(-R(\tilde{t}))}.$$

Then the recognizability of representations can be estimated based on the outputs of  $\hat{\xi}$  via Eq. (4). We find that the recognizability of representations can be used as an indicator of the generalizability and achieve comparable performance to the  $f$ -CMI bound, which will be discussed in the next section.

## 4 Experiments

In this section, we demonstrate the effectiveness of RIB as a training objective to regularize the DNNs and verify the capability of recognizability critic to estimate generalization gap. Code is available at <https://github.com/lvyilin/RecogIB>.

### 4.1 Experimental Setup

**Datasets.** Our experiments mainly conduct on three widely-used datasets: Fashion-MNIST [Xiao *et al.*, 2017], SVHN [Netzer *et al.*, 2011] and CIFAR10 [Krizhevsky and Hinton, 2009]. We also give the results on MNIST and STL10 [Coates *et al.*, 2011] in Appendix C. We adopt the data setting of [Harutyunyan *et al.*, 2021] to enable the calculation and comparison of  $f$ -CMI bound. Concretely, each experiment is conducted on  $k_1$  draws of the original training set (to form a supersample  $\tilde{z}$ ) and  $k_2$  draws of the train-val-split (i.e.,  $k_1 k_2$  runs in total), which correspond to the randomness of  $\tilde{Z}$  and  $T$ , respectively. In all our experiments,  $k_1 = k_2 = 5$ . Unless otherwise stated, the validation set is used as the ghost set. To demonstrate the effectiveness across different data scales, we perform sub-sampling on the training set with different sizes  $n \in \{1250, 5000, 20000\}$ .

**Implementation details.** We use a DNN model composed of a 4-layer CNN (128-128-256-1024) and a 2-layer MLP (1024-512) as the encoder, and use a 4-layer MLP (1024-1024-1024-1) as the recognizability critic. We train the learning model using Adam optimizer [Kingma and Ba, 2015] with betas of (0.9, 0.999) and train the recognizability critic using SGD with momentum of 0.9 as it is more stable in practice. All learning rates are set to 0.001, and the models are trained 100 epochs using the cosine annealing learning rate scheme with a batch size of 128. The trade-off parameter  $\beta$  is selected from  $\{10^{-1}, 10^0, 10^1, 10^2\}$  according to the desired regularization strength as discussed later. All the experiments are implemented with PyTorch and performed on eight NVIDIA RTX A4000 GPUs. More experimental details are provided in Appendix B.

### 4.2 Regularization Effects of RIB

We report the average performance of RIB over  $k_1 k_2$  draws to show its regularization effects on different data scales. We use cross entropy (CE) loss as the baseline. Furthermore, we compare with two common regularization methods: L2 normalization and dropout [Srivastava *et al.*, 2014]. The strength of the L2 normalization is set to 1e-4 and the dropout rate is set to 0.1. We also compare with four popular IB counterparts: VIB [Alemi *et al.*, 2017], NIB [Kolchinsky *et al.*, 2019], DIB [Yu *et al.*, 2021], and PIB [Wang *et al.*, 2022]. All methods impose regularization on the representations, except for PIB which is on the weights. Their regularization strength is determined by the best result of  $\{10^{-5}, 10^{-4}, \dots, 10^1\}$ . Table 1 shows that RIB consistently outperforms the compared methods across all datasets and all training set sizes. This demonstrates that regularizing the recognizability of representations as well as reducing the recognizable information contributes significantly to improving the generalization of the model. Another point worth mentioning is that the performance of PIB does not outperform VIB and its variants. We consider that this is because the PIB computes the gradient for the overall data only once per epoch, in order to maintain an affordable computational complexity, while representation-based regularization methods compute the gradient for each iteration efficiently, thus provide more fine-grained information for model training. Furthermore, we credit the improve-

Method	Fashion			SVHN			CIFAR10		
	1250	5000	20000	1250	5000	20000	1250	5000	20000
CE	18.61±0.87	14.57±0.47	10.18±0.34	30.03±1.04	19.51±0.52	11.77±0.33	61.51±1.13	50.22±0.81	32.81±0.70
L2	18.54±0.66	14.81±0.39	10.57±0.50	29.60±1.00	19.30±0.73	11.19±0.22	61.62±0.97	49.89±0.93	32.41±1.17
Dropout	18.58±0.57	14.58±0.51	10.20±0.40	29.68±0.94	19.40±0.57	11.82±0.42	61.53±0.95	50.02±0.78	32.69±0.47
VIB	17.57±0.44	13.86±0.35	9.64±0.24	40.54±5.75	18.93±0.61	11.28±0.23	59.77±0.78	49.14±0.60	31.82±0.39
NIB	18.30±0.42	14.34±0.29	9.94±0.21	29.34±0.71	19.44±0.38	11.59±0.24	61.00±0.77	50.01±0.57	32.32±0.41
DIB	18.11±0.37	13.99±0.34	9.74±0.27	29.95±0.93	19.37±0.62	11.54±0.23	61.02±0.70	50.09±0.66	32.36±0.49
PIB	18.44±0.52	14.33±0.29	10.00±0.23	29.40±0.78	19.40±0.61	11.60±0.27	61.19±0.82	50.09±0.79	32.48±0.31
RIB	<b>17.05±0.41</b>	<b>13.57±0.28</b>	<b>9.32±0.25</b>	<b>27.72±1.17</b>	<b>17.37±0.52</b>	<b>10.80±0.27</b>	<b>59.16±1.17</b>	<b>47.57±0.60</b>	<b>31.06±0.38</b>

Table 1: Comparison of the mean test error (%) on three datasets with different training set sizes.




Figure 2: Mean test risk curves on CIFAR10 with training set sizes of 1250 and 20000. “RIB-adv” represents optimizing the RIB objective with adversarial training.

Source	1250	5000	20000
SVHN	61.26±1.20	51.59±0.97	37.13±0.94
CIFAR100	60.31±0.91	49.20±0.78	33.27±0.61
CIFAR10	<b>59.16±1.17</b>	<b>47.57±0.60</b>	<b>31.06±0.38</b>
Baseline	61.51±1.13	50.22±0.81	32.81±0.70

Table 2: Comparison of the mean test error (%) using CIFAR10 as the training data and constructing ghosts with different sources.

ment of RIB over VIB and its variants to the use of the ghost set, which enables the computation of recognizable information, and the reduction of recognizability of representations to facilitate generalization.

**Effectiveness of the surrogate loss.** We first examine whether optimizing the Bregman divergence of Eq. (8) is better than optimizing Eq. (7) by adversarial training method using Jensen-Shannon divergence [Nowozin *et al.*, 2016]. The risk curves are illustrated in Fig. 2. We find that: *i*) training RIB in an adversarial manner may produce unstable gradients when the data is scarce thus making the model difficult to converge; *ii*) with more data, “RIB-adv” can also regularize the model and perform better than the CE baseline; and *iii*) the RIB optimized by our surrogate loss consistently outperforms the baselines at different data scales, which demonstrates its effectiveness on regularizing the DNNs.

**Results of ghosts with different sources.** Although the construction of the ghost set is supposed to be performed on a supersample drawn from a single data domain, it is intriguing

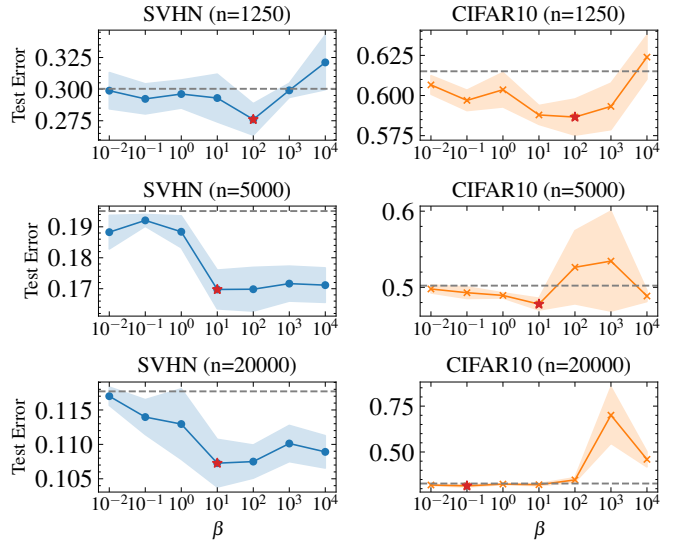


Figure 3: Impact of  $\beta$  on mean test risk with different training set sizes. The horizontal dashed line indicates the performance on CE baseline. The star marker indicates the best result.

to investigate whether it can still achieve similar performance with different data sources. We use the CIFAR10 as the training set and construct the ghost from three sources: *i*) SVHN, distinct from CIFAR10; *ii*) CIFAR100 [Krizhevsky and Hinton, 2009], similar to CIFAR10 but with more categories and *iii*) CIFAR10 for comparison. Table 2 shows that *i*) the more similar the sources are to the training set, the more effective the regularization will be. This is reasonable because recognizing samples from other domains may be too simple thus not provide useful information. More similar sources can provide more fine-grained recognizable information that is of more benefit to the model. *ii*) Nevertheless, even data from different sources, such as SVHN, can still provide valuable information to improve generalization when training data is scarce (e.g.,  $n = 1250$ ).

**Impacts of  $\beta$ .** To investigate the impacts of the regularization term of the RIB objective, we train models using the RIB objective with various  $\beta$ . The results on SVHN and CIFAR10 are illustrated in Fig. 3. We can observe that: *i*) When training data is scarce, the model obtains the best performance at




Figure 4: Comparison of expected generalization gap, mean test error,  $f$ -CMI bound and recognizability of representations on three datasets with varying training set sizes.




Figure 5: (a)(b)(c) Visualization of the achievable region corresponding to the recognizability of representations on three datasets with varying training set sizes. Each polygon represents one run of the experiment. Smaller polygons indicate less recognizability. (d) Dynamics of recognizability on CIFAR10 with size  $n = 1250$ .

larger  $\beta$ . This is expected since the model tends to over-fit to the data in this case. *ii*) SVHN is more tolerant of large  $\beta$  compared to CIFAR10. When increasing the  $\beta$  to more than  $10^3$ , the model still outperforms the CE baseline in most cases, while too large  $\beta$  may lead to instability in CIFAR10 training. We consider this because the features of SVHN are easier to learn than those of CIFAR10 and are more likely to be over-fitted to irrelevant features, thus regularization brings greater improvement. This provides us with some indications for choosing the best  $\beta$  based on the characteristics of the dataset.

### 4.3 Estimation of Generalization Gap

To verify the recognizability of representations can be used to indicate the generalization performance, we conduct experiments on different training set sizes  $n \in \{1250, 5000, 20000\}$  with the smaller dataset being a subset of the larger dataset. We first train the models on different dataset sizes respectively, and then train the recognizability critics using the representations obtained from the trained models on the smallest data subset respectively. This allows us to evaluate the generalization to the unseen samples outside the smallest subset. Each experiment is repeated  $k_1 k_2$  times to evaluate the  $f$ -CMI bound. As illustrated in Fig. 4, we find that *i*) both the recognizability of representations and the  $f$ -CMI bound reflect the generalization to the unseen data, and the recognizability of representations even obtains better estimates in some cases; *ii*) since the recognizability of representations is a parameterized approximation rather than an upper bound, the risk may sometimes be underestimated as in Fig. 4b. We also visualize the achievable region corresponding to the recognizability of representations, which is obtained by plotting the ROC curve

of the score function and its central symmetry about the point  $(1/2, 1/2)$ . The symmetry of the achievable region is due to the fact that we can construct an opposite test when any test chooses  $H_0$  it chooses  $H_1$  and vice versa. The visualization results are illustrated in Fig. 5. We find that the better the generalization performance, the smaller the area of achievable regions, which validates the plausibility of our formulation. Since CIFAR10 features are more difficult to learn thus easier to over-fit, it can be seen that a larger area is covered than the other two datasets.

**Recognizability Dynamics.** To investigate how the recognizability changes during training, we evaluate the recognizability every 5 epochs and plot the dynamics of recognizability as the number of training epochs changes in Fig. 5d. We observe that recognizability can track the generalization gap as expected and eventually converge to around the  $f$ -CMI bound. See Appendix C for more results.

## 5 Conclusions and Future Work

In this paper, we establish the connection between the recognizability of representations and the  $f$ -CMI bound, and propose the RIB to regularize the recognizability of representations through Bregman divergence. We conducted extensive experiments on several commonly used datasets to demonstrate the effectiveness of RIB in regularizing the model and estimating the generalization gap. Future work can investigate whether RIB is also effective on other tasks such as few-shot learning, out-of-distribution detection, domain generalization, etc.

## Acknowledgments

This work was partly supported by the National Natural Science Foundation of China under Grant 62176020; the National Key Research and Development Program (2020AAA0106800); the Beijing Natural Science Foundation under Grant L211016; the Fundamental Research Funds for the Central Universities (2019JBZ110); Chinese Academy of Sciences (OEIP-O-202004); Grant ITF MHP/038/20, Grant CRF 8730063, Grant RGC 14300219, 14302920.

## References

- [Aleml *et al.*, 2017] Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [Banerjee and Montúfar, 2021] Pradeep Kr. Banerjee and Guido Montúfar. Information complexity and generalization bounds. In *IEEE International Symposium on Information Theory, ISIT 2021, Melbourne, Australia, July 12-20, 2021*, pages 676–681. IEEE, 2021.
- [Bu *et al.*, 2020] Yuheng Bu, Shaofeng Zou, and Venugopal V. Veeravalli. Tightening mutual information-based bounds on generalization error. *IEEE J. Sel. Areas Inf. Theory*, 1(1):121–130, 2020.
- [Catoni, 2007] Olivier Catoni. Pac-bayesian supervised classification: The thermodynamics of statistical learning. *IMS Lecture Notes Monograph Series*, 56:1–163, 2007. arXiv:0712.0248 [stat].
- [Coates *et al.*, 2011] Adam Coates, Andrew Y. Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In Geoffrey J. Gordon, David B. Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011*, volume 15, pages 215–223. JMLR.org, 2011.
- [Du *et al.*, 2020] Ying-Jun Du, Jun Xu, Huan Xiong, Qiang Qiu, Xiantong Zhen, Cees G. M. Snoek, and Ling Shao. Learning to learn with variational information bottleneck for domain generalization. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part X*, volume 12355, pages 200–216. Springer, 2020.
- [Dubois *et al.*, 2020] Yann Dubois, Douwe Kiela, David J. Schwab, and Ramakrishna Vedantam. Learning optimal representations with the decodable information bottleneck. In *NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [Dziugaite and Roy, 2017] Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*. AUAI Press, 2017.
- [Gálvez *et al.*, 2020] Borja Rodríguez Gálvez, Ragnar Thobaben, and Mikael Skoglund. The convex information bottleneck lagrangian. *Entropy*, 22(1):98, 2020.
- [Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 2672–2680, 2014.
- [Haghifam *et al.*, 2020] Mahdi Haghifam, Jeffrey Negrea, Ashish Khisti, Daniel M. Roy, and Gintare Karolina Dziugaite. Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms. In *NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [Harutyunyan *et al.*, 2021] Hrayr Harutyunyan, Maxim Raginsky, Greg Ver Steeg, and Aram Galstyan. Information-theoretic generalization bounds for black-box learning algorithms. In *NeurIPS 2021, December 6-14, 2021, virtual*, pages 24670–24682, 2021.
- [Hjelm *et al.*, 2019] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [Kolchinsky *et al.*, 2019] Artemy Kolchinsky, Brendan D. Tracey, and David H. Wolpert. Nonlinear information bottleneck. *Entropy*, 21(12):1181, 2019.
- [Krizhevsky and Hinton, 2009] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009.
- [Lai *et al.*, 2021] Qiuxia Lai, Yu Li, Ailing Zeng, Minhao Liu, Hanqiu Sun, and Qiang Xu. Information bottleneck approach to spatial attention learning. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 779–785. ijcai.org, 2021.
- [Levy, 2008] Bernard C. Levy. *Binary and M-ary Hypothesis Testing*, page 1–57. Springer US, Boston, MA, 2008.
- [Li *et al.*, 2022] Bo Li, Yifei Shen, Yezhen Wang, Wenzhen Zhu, Colorado Reed, Dongsheng Li, Kurt Keutzer, and Han Zhao. Invariant information bottleneck for domain



- generalization. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, Virtual Event, February 22 - March 1, 2022*, pages 7399–7407. AAAI Press, 2022.
- [Negrea *et al.*, 2019] Jeffrey Negrea, Mahdi Haghifam, Gintare Karolina Dziugaite, Ashish Khisti, and Daniel M. Roy. Information-theoretic generalization bounds for SGLD via data-dependent estimates. In *NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 11013–11023, 2019.
- [Netzer *et al.*, 2011] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- [Neyman *et al.*, 1933] Jerzy Neyman, Egon Sharpe Pearson, and Karl Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694–706):289–337, Feb 1933.
- [Nowozin *et al.*, 2016] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *NeurIPS 2016, December 5-10, 2016, Barcelona, Spain*, pages 271–279, 2016.
- [Pan *et al.*, 2021] Ziqi Pan, Li Niu, Jianfu Zhang, and Liqing Zhang. Disentangled information bottleneck. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Event, February 2-9, 2021*, pages 9285–9293. AAAI Press, 2021.
- [Pensia *et al.*, 2018] Ankit Pensia, Varun S. Jog, and Po-Ling Loh. Generalization error bounds for noisy, iterative algorithms. In *2018 IEEE International Symposium on Information Theory, ISIT 2018, Vail, CO, USA, June 17-22, 2018*, pages 546–550. IEEE, 2018.
- [Poole *et al.*, 2019] Ben Poole, Sherjil Ozair, Aaron van den Oord, Alexander A. Alemi, and George Tucker. On variational bounds of mutual information. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97, pages 5171–5180. PMLR, 2019.
- [Rodriguez Galvez, 2019] Borja Rodriguez Galvez. The information bottleneck : Connections to other problems, learning and exploration of the ib curve. Master’s thesis, KTH, School of Electrical Engineering and Computer Science (EECS), 2019.
- [Saxe *et al.*, 2018] Andrew M. Saxe, Yamini Bansal, Joel Dapello, Madhu Advani, Artemy Kolchinsky, Brendan D. Tracey, and David D. Cox. On the information bottleneck theory of deep learning. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.
- [Shamir *et al.*, 2010] Ohad Shamir, Sivan Sabato, and Naftali Tishby. Learning and generalization with the information bottleneck. *Theor. Comput. Sci.*, 411(29-30):2696–2711, 2010.
- [Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, 2014.
- [Steinke and Zakyntinou, 2020] Thomas Steinke and Lydia Zakyntinou. Reasoning about generalization via conditional mutual information. In *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125, pages 3437–3452. PMLR, 2020.
- [Sugiyama *et al.*, 2012] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. Density-ratio matching under the bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64(5):1009–1044, Oct 2012.
- [Tishby *et al.*, 2000] Naftali Tishby, Fernando C. N. Pereira, and William Bialek. The information bottleneck method. *CoRR*, physics/0004057, 2000.
- [Vera *et al.*, 2018] Matías Vera, Pablo Piantanida, and Leonardo Rey Vega. The role of the information bottleneck in representation learning. In *2018 IEEE International Symposium on Information Theory, ISIT 2018, Vail, CO, USA, June 17-22, 2018*, pages 1580–1584. IEEE, 2018.
- [Wang *et al.*, 2022] Zifeng Wang, Shao-Lun Huang, Erkan Engin Kuruoglu, Jimeng Sun, Xi Chen, and Yefeng Zheng. PAC-bayes information bottleneck. In *International Conference on Learning Representations*, 2022.
- [Xiao *et al.*, 2017] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017.
- [Xu and Raginsky, 2017] Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *NeurIPS 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2524–2533, 2017.
- [Yu *et al.*, 2021] Xi Yu, Shujian Yu, and José C. Príncipe. Deep deterministic information bottleneck with matrix-based entropy functional. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pages 3160–3164. IEEE, 2021.
- [Zhou *et al.*, 2022] Ruida Zhou, Chao Tian, and Tie Liu. Individually conditional individual mutual information bound on generalization error. *IEEE Trans. Inf. Theory*, 68(5):3304–3316, 2022.