

SELF-SUPERVISED REPRESENTATION LEARNING WITH RELATIVE PREDICTIVE CODING

Yao-Hung Hubert Tsai¹, Martin Q. Ma¹, Muqiao Yang¹,
Han Zhao^{2,3}, Louis-Philippe Morency¹, Ruslan Salakhutdinov¹

¹Carnegie Mellon University, ²D.E. Shaw & Co., ³ University of Illinois at Urbana-Champaign

ABSTRACT

This paper introduces Relative Predictive Coding (RPC), a new contrastive representation learning objective that maintains a good balance among training stability, minibatch size sensitivity, and downstream task performance. The key to the success of RPC is two-fold. First, RPC introduces the relative parameters to regularize the objective for boundedness and low variance. Second, RPC contains no logarithm and exponential score functions, which are the main cause of training instability in prior contrastive objectives. We empirically verify the effectiveness of RPC on benchmark vision and speech self-supervised learning tasks. Lastly, we relate RPC with mutual information (MI) estimation, showing RPC can be used to estimate MI with low variance¹.

1 INTRODUCTION

Unsupervised learning has drawn tremendous attention recently because it can extract rich representations without label supervision. Self-supervised learning, a subset of unsupervised learning, learns representations by allowing the data to provide supervision (Devlin et al., 2018). Among its mainstream strategies, self-supervised contrastive learning has been successful in visual object recognition (He et al., 2020; Tian et al., 2019; Chen et al., 2020c), speech recognition (Oord et al., 2018; Rivière et al., 2020), language modeling (Kong et al., 2019), graph representation learning (Velickovic et al., 2019) and reinforcement learning (Kipf et al., 2019). The idea of self-supervised contrastive learning is to learn latent representations such that related instances (e.g., patches from the same image; defined as *positive* pairs) will have representations within close distance, while unrelated instances (e.g., patches from two different images; defined as *negative* pairs) will have distant representations (Arora et al., 2019).

Prior work has formulated the contrastive learning objectives as maximizing the divergence between the distribution of related and unrelated instances. In this regard, different divergence measurement often leads to different loss function design. For example, variational mutual information (MI) estimation (Poole et al., 2019) inspires Contrastive Predictive Coding (CPC) (Oord et al., 2018). Note that MI is also the KL-divergence between the distributions of related and unrelated instances (Cover & Thomas, 2012). While the choices of the contrastive learning objectives are abundant (Hjelm et al., 2018; Poole et al., 2019; Ozair et al., 2019), we point out that there are three challenges faced by existing methods.

The first challenge is the training stability, where an unstable training process with high variance may be problematic. For example, Hjelm et al. (2018); Tschannen et al. (2019); Tsai et al. (2020b) show that the contrastive objectives with large variance cause numerical issues and have a poor downstream performance with their learned representations. The second challenge is the sensitivity to minibatch size, where the objectives requiring a huge minibatch size may restrict their practical usage. For instance, SimCLRv2 (Chen et al., 2020c) utilizes CPC as its contrastive objective and reaches state-of-the-art performances on multiple self-supervised and semi-supervised benchmarks. Nonetheless, the objective is trained with a minibatch size of 8,192, and this scale of training requires enormous computational power. The third challenge is the downstream task performance, which is the one that we would like to emphasize the most. For this reason, in most cases, CPC

¹Project page: https://github.com/martinmamql/relative_predictive_coding

Table 1: Different contrastive learning objectives, grouped by measurements of distribution divergence. P_{XY} represents the distribution of related samples (positively-paired), and $P_X P_Y$ represents the distribution of unrelated samples (negatively-paired). $f(x; y) \geq F$ for F being any class of functions $f : X \times Y \rightarrow \mathbb{R}$. Compared to J_{CPC} and J_{RPC} , we empirically find J_{WPC} performs worse on complex real-world image datasets spanning CIFAR-10/-100 (Krizhevsky et al., 2009) and ImageNet (Russakovsky et al., 2015).

Objective	Good Training Stability	Lower Minibatch Size Sensitivity	Good Downstream Performance
relating to KL-divergence between P_{XY} and $P_X P_Y$: J_{DV} (Donsker & Varadhan, 1975), J_{NWJ} (Nguyen et al., 2010), and J_{CPC} (Oord et al., 2018)			
$J_{DV}(X, Y) := \sup_{f \in F} \mathbb{E}_{P_{XY}} [f(x, y)] - \log(\mathbb{E}_{P_X P_Y} [e^{f(x, y)}])$	✗	✓	✗
$J_{NWJ}(X, Y) := \sup_{f \in F} \mathbb{E}_{P_{XY}} [f(x, y)] - \mathbb{E}_{P_X P_Y} [e^{f(x, y)}]$	✗	✓	✗
$J_{CPC}(X, Y) := \sup_{f \in F} \mathbb{E}_{(x, y) \sim P_{XY}} [f(x, y)] - \frac{1}{N} \sum_{j=1}^N \mathbb{E}_{y \sim P_Y} [e^{f(x, y_j)}]$	✓	✗	✓
relating to JS-divergence between P_{XY} and $P_X P_Y$: J_{JS} (Nowozin et al., 2016)			
$J_{JS}(X, Y) := \sup_{f \in F} \mathbb{E}_{P_{XY}} [\log(1 + e^{f(x, y)})] - \mathbb{E}_{P_X P_Y} [\log(1 + e^{f(x, y)})]$	✓	✓	✗
relating to Wasserstein-divergence between P_{XY} and $P_X P_Y$: J_{WPC} (Ozair et al., 2019), with F_L denoting the space of 1-Lipschitz functions			
$J_{WPC}(X, Y) := \sup_{f \in F_L} \mathbb{E}_{(x, y) \sim P_{XY}} [f(x, y)] - \frac{1}{N} \sum_{j=1}^N \mathbb{E}_{y \sim P_Y} [e^{f(x, y_j)}]$	✓	✓	✗
relating to χ^2 -divergence between P_{XY} and $P_X P_Y$: J_{RPC} (ours)			
$J_{RPC}(X, Y) := \sup_{f \in F} \mathbb{E}_{P_{XY}} [f(x, y)] - \frac{\alpha}{\beta} \mathbb{E}_{P_X P_Y} [f(x, y)] - \frac{\beta}{\alpha} \mathbb{E}_{P_X P_Y} [f^2(x, y)]$	✓	✓	✓

is the objective that we would adopt for contrastive representation learning, due to its favorable performance in downstream tasks (Tschannen et al., 2019; Baevski et al., 2020).

This paper presents a new contrastive representation learning objective: the Relative Predictive Coding (RPC), which attempts to achieve a good balance among these three challenges: training stability, sensitivity to minibatch size, and downstream task performance. At the core of RPC is the *relative parameters*, which are used to regularize RPC for its boundedness and low variance. From a modeling perspective, the relative parameters act as a ℓ_2 regularization for RPC. From a statistical perspective, the relative parameters prevent RPC from growing to extreme values, as well as upper bound its variance. In addition to the relative parameters, RPC contains no logarithm and exponential, which are the main cause of the training instability for prior contrastive learning objectives (Song & Ermon, 2019).

To empirically verify the effectiveness of RPC, we consider benchmark self-supervised representation learning tasks, including visual object classification on CIFAR-10/-100 (Krizhevsky et al., 2009), STL-10 (Coates et al., 2011), and ImageNet (Russakovsky et al., 2015) and speech recognition on LibriSpeech (Panayotov et al., 2015). Comparing RPC to prior contrastive learning objectives, we observe a lower variance during training, a lower minibatch size sensitivity, and consistent performance improvement. Lastly, we also relate RPC with MI estimation, empirically showing that RPC can estimate MI with low variance.

2 PROPOSED METHOD

This paper presents a new contrastive representation learning objective - the Relative Predictive Coding (RPC). At a high level, RPC 1) introduces the relative parameters to regularize the objective for boundedness and low variance; and 2) achieves a good balance among the three challenges in the contrastive representation learning objectives: training stability, sensitivity to minibatch size, and downstream task performance. We begin by describing prior contrastive objectives along with their limitations on the three challenges in Section 2.1. Then, we detail our presented objective and its modeling benefits in Section 2.2. An overview of different contrastive learning objectives is provided in Table 1. We defer all the proofs in Appendix.

Notation We use an uppercase letter to denote a random variable (e.g., X), a lower case letter to denote the outcome of this random variable (e.g., x), and a calligraphy letter to denote the sample space of this random variable (e.g., \mathcal{X}). Next, if the samples (x, y) are related (or positively-paired), we refer $(x, y) \sim P_{XY}$ with P_{XY} being the joint distribution of $X \times Y$. If the samples (x, y) are unrelated (negatively-paired), we refer $(x, y) \sim P_X P_Y$ with $P_X P_Y$ being the product of marginal distributions over $X \times Y$. Last, we define $f \geq F$ for F being any class of functions $f : X \times Y \rightarrow \mathbb{R}$.

2.1 PRELIMINARY

Contrastive representation learning encourages the *contrastiveness* between the positive and the negative pairs of the representations from the related data X and Y . Specifically, when sampling a pair

of representation $(x; y)$ from their joint distribution $(x; y) \sim P_{XY}$, this pair is defined as a positive pair; when sampling from the product of marginals $(x; y) \sim P_X P_Y$, this pair is defined as a negative pair. Then, Tsai et al. (2020b) formalizes this idea such that the contrastiveness of the representations can be measured by the divergence between P_{XY} and $P_X P_Y$, where higher divergence suggests better contrastiveness. To better understand prior contrastive learning objectives, we categorize them in terms of different divergence measurements between P_{XY} and $P_X P_Y$, with their detailed objectives presented in Table 1.

We instantiate the discussion using Contrastive Predictive Coding (Oord et al., 2018), which is a lower bound of $D_{KL}(P_{XY} \parallel P_X P_Y)$ with D_{KL} referring to the KL-divergence:

$$J_{CPC}(X; Y) := \sup_{f \in \mathcal{F}} E_{(x; y_1) \sim P_{XY}; (y_2; y_3) \sim P_Y} \log \frac{e^{f(x; y_1)}}{\frac{1}{N} \sum_{j=1}^N e^{f(x; y_j)}} \quad (1)$$

Then, Oord et al. (2018) presents to maximize $J_{CPC}(X; Y)$, so that the learned representations X and Y have high contrastiveness. We note that J_{CPC} has been commonly used in many recent self-supervised representation learning frameworks (He et al., 2020; Chen et al., 2020b), where they constrain the function to be $f(x; y) = \text{cosine}(x; y)$ with $\text{cosine}()$ being cosine similarity. Under this function design, maximizing J_{CPC} leads the representations of related pairs to be close and representations of unrelated pairs to be distant.

The category of modeling $D_{KL}(P_{XY} \parallel P_X P_Y)$ also includes the Donsker-Varadhan objective (J_{DV} (Donsker & Varadhan, 1975; Belghazi et al., 2018)) and the Nguyen-Wainwright-Jordan objective (J_{NWJ} (Nguyen et al., 2010; Belghazi et al., 2018)), where Belghazi et al. (2018); Tsai et al. (2020b) show that $J_{DV}(X; Y) = J_{NWJ}(X; Y) = D_{KL}(P_{XY} \parallel P_X P_Y)$. The other divergence measurements considered in prior work are $D_{JS}(P_{XY} \parallel P_X P_Y)$ (with D_{JS} referring to the Jensen-Shannon divergence) and $D_{Wass}(P_{XY} \parallel P_X P_Y)$ (with D_{Wass} referring to the Wasserstein-divergence). The instance of modeling $D_{JS}(P_{XY} \parallel P_X P_Y)$ is the Jensen-Shannon f-GAN objective J_{JS} (Nowozin et al., 2016; Hjelm et al., 2018) where $J_{JS}(X; Y) = 2 D_{JS}(P_{XY} \parallel P_X P_Y) \log 2$.² The instance of modeling $D_{Wass}(P_{XY} \parallel P_X P_Y)$ is the Wasserstein Predictive Coding J_{WPC} (Ozair et al., 2019), where $J_{WPC}(X; Y)$ modifies $J_{CPC}(X; Y)$ objective (equation 1) by searching the function from \mathcal{F} to F_L . F_L denotes any class of 1-Lipschitz continuous functions from $(X \times Y)$ to \mathbb{R} , and thus $F_L \subseteq \mathcal{F}$. Ozair et al. (2019) shows that $J_{WPC}(X; Y)$ is the lower bound of both $D_{KL}(P_{XY} \parallel P_X P_Y)$ and $D_{Wass}(P_{XY} \parallel P_X P_Y)$. See Table 1 for all the equations. To conclude, the contrastive representation learning objectives are unsupervised representation learning methods that maximize the distribution divergence between P_{XY} and $P_X P_Y$. The learned representations cause high contrastiveness, and recent work (Arora et al., 2019; Tsai et al., 2020a) theoretically show that highly-contrastive representations could improve the performance on downstream tasks.

After discussing prior contrastive representation learning objectives, we point out three challenges in their practical deployments: training stability, sensitivity to minibatch training size, and downstream task performance. In particular, the three challenges can hardly be handled well at the same time, where we highlight the conclusions in Table Training Stability: The training stability highly relates to the variance of the objectives, where Song & Ermon (2019) shows that J_{DV} and J_{NWJ} exhibit inevitable high variance due to their inclusion of exponential function. As pointed out by Tsai et al. (2020b), J_{CPC} , J_{WPC} , and J_{JS} have better training stability because J_{CPC} and J_{WPC} can be realized as a multi-class classification task and can be realized as a binary classification task. The cross-entropy loss adopted in J_{CPC} , J_{WPC} , and J_{JS} is highly-optimized and stable in existing optimization package (Abadi et al., 2016; Paszke et al., 2019). Sensitivity to minibatch training size: Among all the prior contrastive representation learning methods, J_{CPC} is known to be sensitive to the minibatch training size (Ozair et al., 2019). Taking a closer look at equation 1, J_{CPC} deploys an instance selection such that should be selected from $(y_1; y_2; \dots; y_N)$, with $(x; y_1) \sim P_{XY}$, $(x; y_j) \sim P_X P_Y$ with N being the minibatch size. Previous work (Poole et al., 2019; Song & Ermon, 2019; Chen et al., 2020b; Caron et al., 2020) showed that J_{CPC} results in a more challenging instance selection and forces to have a better contrastiveness of (related instance for x) against $(y_j)_{j=2}^N$ (unrelated instance for x). J_{DV} , J_{NWJ} , and J_{JS} do not consider

² $J_{JS}(X; Y)$ achieves its supreme value when $f(x; y) = \log(p(x; y) / (p(x)p(y)))$ (Tsai et al., 2020b). Plug in $f(x; y)$ into $J_{JS}(X; Y)$, we can conclude $J_{JS}(X; Y) = 2(D_{JS}(P_{XY} \parallel P_X P_Y) \log 2)$.

the instance selection, and J_{WPC} reduces the minibatch training size sensitivity by enforcing 1-Lipschitz constraint. Downstream Task Performance The downstream task performance is what we care the most among all the three challenges. J_{CPC} has been the most popular objective as it manifests superior performance over the other alternatives (Tschannen et al., 2019; Tsai et al., 2020b;a). We note that although J_{WPC} shows better performance on Omniglot (Lake et al., 2015) and CelebA (Liu et al., 2015) datasets, we empirically find it not generalizing well to CIFAR-10/100 (Krizhevsky et al., 2009) and ImageNet (Russakovsky et al., 2015).

2.2 RELATIVE PREDICTIVE CODING

In this paper, we present Relative Predictive Coding (RPC), which achieves a good balance among the three challenges mentioned above:

$$J_{RPC}(X; Y) := \sup_{f \in \mathcal{F}} E_{P_{XY}} [f(x; y)] - E_{P_X P_Y} [f(x; y)] - \frac{1}{2} E_{P_{XY}} \|f\|^2(x; y) - \frac{1}{2} E_{P_X P_Y} \|f\|^2(x; y); \quad (2)$$

where $\lambda > 0$, $\mu > 0$, $\nu > 0$ are hyper-parameters and we define them relative parameters. Intuitively, J_{RPC} contains no logarithm or exponential, potentially preventing unstable training due to numerical issues. Now, we discuss the roles of λ . At a first glance, λ acts to discourage the scores of P_{XY} and $P_X P_Y$ from being close, and μ acts as a regularization coefficient to stop f from becoming large. For a deeper analysis, the relative parameters act to regularize our objective for boundedness and low variance. To show this claim, we first present the following lemma:

Lemma 1 (Optimal Solution for J_{RPC}) Let $r(x; y) = \frac{p(x; y)}{p(x)p(y)}$ be the density ratio. J_{RPC} has the optimal solution $f^*(x; y) = \frac{r(x; y) - 1}{r(x; y) + 1} := r_{\lambda; \mu; \nu}(x; y)$ with $\lambda = \frac{1}{2}$.

Lemma 1 suggests that J_{RPC} achieves its supreme value at the ratio $r_{\lambda; \mu; \nu}(x; y)$ indexed by the relative parameters $\lambda; \mu; \nu$ (i.e., we term $r_{\lambda; \mu; \nu}(x; y)$ as the relative density ratio). We note that $r_{\lambda; \mu; \nu}(x; y)$ is an increasing function w.r.t. $r(x; y)$ and is nicely bounded even when $r(x; y)$ is large. We will now show that the boundedness of $r_{\lambda; \mu; \nu}$ suggests the empirical estimation $J_{RPC}^{m; n}$ has boundedness and low variance. In particular, let $y_j^0, g_{j=1}^m$ be m samples drawn uniformly at random from P_X and $x_j^0, y_j^0, g_{j=1}^m$ be m samples drawn uniformly at random from P_Y . Then, we use neural networks to empirically estimate J_{RPC} as $J_{RPC}^{m; n}$:

Definition 1 ($J_{RPC}^{m; n}$, empirical estimation of J_{RPC}) We parametrize \mathcal{F} via a family of neural networks $\mathcal{F} := \{f : \mathbb{R}^d \rightarrow \mathbb{R}\}$ where $d \geq 2$ and \mathbb{R}^d is compact. Then, $J_{RPC}^{m; n} = \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n f(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) - \frac{1}{n} \sum_{i=1}^n \frac{1}{2} f^2(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m \frac{1}{2} f^2(x_j^0; y_j^0)$.

Proposition 1 (Boundedness of $J_{RPC}^{m; n}$, informal) $0 \leq J_{RPC} \leq \frac{1}{2} + \frac{1}{2}$. Then, with probability at least $1 - \delta$, $J_{RPC} \leq J_{RPC}^{m; n} \leq O\left(\frac{d + \log(1/\delta)}{n^0}\right)$; where $n^0 = \min\{n, m\}$.

Proposition 2 (Variance of $J_{RPC}^{m; n}$, informal) There exist universal constants c_1 and c_2 that depend only on $\lambda; \mu; \nu$, such that $\text{Var}[J_{RPC}^{m; n}] = O\left(\frac{c_1}{n} + \frac{c_2}{m}\right)$.

From the two propositions, when n and m are large, i.e., the sample sizes are large, $J_{RPC}^{m; n}$ is bounded, and its variance vanishes. First, the boundedness of $J_{RPC}^{m; n}$ suggests $J_{RPC}^{m; n}$ will not grow to extremely large or small values. Prior contrastive learning objectives with good training stability (e.g., $J_{CPC} / J_{JS} / J_{WPC}$) also have the boundedness of their objective values. For instance, the empirical estimation of J_{CPC} is less than $\log N$ (equation 1) (Poole et al., 2019). Nevertheless, J_{CPC} often performs the best only when minibatch size is large, and empirical performances of J_{WPC} are not as competitive as J_{CPC} . Second, the upper bound of the variance implies the training of $J_{RPC}^{m; n}$ can be stable, and in practice we observe a much smaller value than the stated upper bound. On the contrary, Song & Ermon (2019) shows that the empirical estimations of J_{NWJ} exhibit inevitable variances that grow exponentially with the size of $(P_{XY} \times P_X \times P_Y)$.

Lastly, similar to prior contrastive learning objective that are related to distribution divergence measurement, we associate J_{RPC} with the Chi-square divergence $\chi^2(P_{XY} \times P_X \times P_Y) =$

$E_{P_X P_Y} [r^2(x; y)] = 1$ (Nielsen & Nock, 2013). The derivations are provided in Appendix. By having $P^0 = \frac{1}{2} P_{XY} + \frac{1}{2} P_X P_Y$ as the mixture distribution dP_{XY} and $P_X P_Y$, we can rewrite $J_{RPC}(X; Y)$ as $J_{RPC}(X; Y) = \frac{1}{2} E_{P^0} [r^2(x; y)]$. Hence, J_{RPC} can be regarded as a generalization of D_2 with the relative parameters α, β , where D_2 can be recovered from J_{RPC} by specializing $\alpha = 0, \beta = 0$ and $\gamma = 1$ (e.g., $D_2 = 2 J_{RPC}(\alpha = 0, \beta = 0, \gamma = 1)$). Note that J_{RPC} may not be a formal divergence measure with arbitrary α, β, γ .

3 EXPERIMENTS

We provide an overview of the experimental section. First, we conduct benchmark self-supervised representation learning tasks spanning visual object classification and speech recognition. This set of experiments are designed to discuss the three challenges of the contrastive representation learning objectives: downstream task performance (Section 3.1), training stability (Section 3.2), and mini-batch size sensitivity (Section 3.3). We also provide an ablation study on the choices of the relative parameters in J_{RPC} (Section 3.4). On these experiments we found that J_{RPC} achieves a lower variance during training, a lower batch size insensitivity, and consistent performance improvement. Second, we relate J_{RPC} with mutual information (MI) estimation (Section 3.5). The connection is that MI is an average statistic of the density ratio, and we have shown that the optimal solution of J_{RPC} is the relative density ratio (see Lemma 1). Thus we could estimate MI using the density ratio transformed from the optimal solution dP_{RPC} . On these two sets of experiments, we fairly compare J_{RPC} with other contrastive learning objectives. Particularly, across different objectives, we fix the network, learning rate, optimizer, and batch size (we use the default configurations suggested by the original implementations from Chen et al. (2020c), Rivet et al. (2020) and Tsai et al. (2020b).) The only difference will be the objective itself. In what follows, we perform the first set of experiments. We defer experimental details in the Appendix.

Datasets. For the visual objective classification, we consider CIFAR-10/-100 (Krizhevsky et al., 2009), STL-10 (Coates et al., 2011), and ImageNet (Russakovsky et al., 2015). CIFAR-10/-100 and ImageNet contain labeled images only, while STL-10 contains labeled and unlabeled images. For the speech recognition, we consider LibriSpeech-100h (Panayotov et al., 2015) dataset, which contains 100 hours of 16kHz English speech from 41 speakers with 41 types of phonemes.

Training and Evaluation Details. For the vision experiments, we follow the setup from SimCLRv2 (Chen et al., 2020c), which considers visual object recognition as its downstream task. For the speech experiments, we follow the setup from prior work (Oord et al., 2018; Rivet et al., 2020), which consider phoneme classification and speaker identification as the downstream tasks. Then, we briefly discuss the training and evaluation details into three modules: 1) related and unrelated data construction, 2) pre-training, and 3) fine-tuning and evaluation. For more details, please refer to Appendix or the original implementations.

Related and Unrelated Data Construction. In the vision experiment, we construct the related images by applying different augmentations on the same image. Hence, $x, y \in P_{XY}$, x and y are the same image with different augmentations. The unrelated images are two randomly selected samples. In the speech experiment, we define the current latent feature (feature at time t) and the future samples (samples at time $t + \Delta$) as related data. In other words, the feature in the latent space should contain information that can be used to infer future time steps. A latent feature and randomly selected samples would be considered as unrelated data.

Pre-training. The pre-training stage refers to the self-supervised training by a contrastive learning objective. Our training objective is defined in Definition 1, where we use neural networks to parametrize the function using the constructed related and unrelated data. Convolutional neural networks are used for vision experiments. Transformers (Vaswani et al., 2017) and LSTMs (Hochreiter & Schmidhuber, 1997) are used for speech experiments.

Fine-tuning and Evaluation. After the pre-training stage, we fix the parameters in the pre-trained networks and add a small fine-tuning network on top of them. Then, we fine-tune this small network with the downstream labels in the data's training split. For the fine-tuning network, both vision and speech experiments consider multi-layer perceptrons. Last, we evaluate the fine-tuned representations on the data's test split. We would like to point out that we do not normalize the hidden representations encoded by the pre-training neural network for loss calculation. This hidden nor-

Table 2: Top-1 accuracy (%) for visual object recognition results. J_{DV} and J_{NWJ} are not reported on ImageNet due to numerical instability. ResNet depth, width and Selective Kernel (SK) configuration for each setting are provided in ResNet Depth+width+SK column. A slight drop of J_{CPC} performance compared to Chen et al. (2020c) is because we only train for 800 epochs rather than 8000 due to the fact that running 800 epochs uninterruptedly on cloud TPU is very expensive. Also, we did not employ a memory buffer (He et al., 2020) to store negative samples. We and we did not employ a memory buffer. We also provide the results from fully supervised models as a comparison (Chen et al., 2020b;c). Fully supervised training performs worse on STL-10 because it does not employ the unlabeled samples in the dataset (Fu et al., 2019).

Dataset	ResNet Depth+Width+SK	Self-supervised						Supervised
		J_{DV}	J_{NWJ}	J_{JS}	J_{WPC}	J_{CPC}	J_{RPC}	
CIFAR-10	18 + 1 + No SK	91.10	90.54	83.55	80.02	91.12	91.46	93.12
CIFAR-10	50 + 1 + No SK	92.23	92.67	87.34	85.93	93.42	93.57	95.70
CIFAR-100	18 + 1 + No SK	77.10	77.27	74.02	72.16	77.36	77.98	79.11
CIFAR-100	50 + 1 + No SK	79.02	78.52	75.31	73.23	79.31	79.89	81.20
STL-10	50 + 1 + No SK	82.25	81.17	79.07	76.50	83.40	84.10	71.40
ImageNet	50 + 1 + SK	-	-	66.21	62.10	73.48	74.43	78.50
ImageNet	152 + 2 + SK	-	-	71.12	69.51	77.80	78.40	80.40

Table 3: Accuracy (%) for LibriSpeech-100h phoneme and speaker classification results. We also provide the results from fully supervised model as a comparison (Oord et al., 2018).

Task Name	Self-supervised				Supervised
	J_{CPC}	J_{DV}	J_{NWJ}	J_{RPC}	
Phoneme classification	64.6	61.27	62.09	69.39	74.6
Speaker classification	97.4	95.36	95.89	97.68	98.5

malization technique is widely applied (Tian et al., 2019; Chen et al., 2020b;c) to stabilize training and increase performance for prior objectives, but we find it unnecessary in

3.1 DOWNSTREAM TASK PERFORMANCES ON VISION AND SPEECH

For the downstream task performance in the vision domain, we test the proposed and other contrastive learning objectives on CIFAR-10/-100 (Krizhevsky et al., 2009), STL-10 (Coates et al., 2011), and ImageNet ILSVRC-2012 (Russakovsky et al., 2015). Here we report the best performances J_{RPC} can get on each dataset (we include experimental details in A.7.) Table 2 shows that the proposed J_{RPC} outperforms other objectives on all datasets. Using on the largest network (ResNet with depth of 152, channel width of 2 and selective kernels), the performance jumps from 77.80% of J_{CPC} to 78.40% of J_{RPC} .

Regarding speech representation learning, the downstream performance for phoneme and speaker classification are shown in Table 3 (we defer experimental details in Appendix A.9.) Compared to J_{CPC} , J_{RPC} improves the phoneme classification results with 10 percent and the speaker classification results with 0.3 percent, which is closer to the fully supervised model. Overall, the proposed J_{RPC} performs better than other unsupervised learning objectives on both phoneme classification and speaker classification tasks.

3.2 TRAINING STABILITY

We provide empirical training stability comparisons by J_{DV} , J_{NWJ} , J_{CPC} and J_{RPC} by plotting the values of the objectives as the training step increases. We apply the four objectives to the SimCLRv2 framework and train on the CIFAR-10 dataset. All setups of training are exactly the same except the objectives. From our experiments, J_{DV} and J_{NWJ} soon explode to NaN and disrupt training (shown as early stopping in Figure 1a; extremely large values are not plotted due to scale constraints). On the other hand, J_{RPC} and J_{CPC} has low variance, and both enjoy stable training. As a result, performances using the representation learned from unstable J_{DV} and J_{NWJ} suffer in downstream task, while representation learned by J_{RPC} and J_{CPC} work much better.

Figure 1: (a) Empirical values of J_{DV} , J_{NWJ} , J_{CPC} and J_{RPC} performing visual object recognition on CIFAR-10. J_{DV} and J_{NWJ} soon explode to NaN values and stop the training (shown as early stopping in the figure), while J_{CPC} and J_{RPC} are more stable. Performance comparison of J_{CPC} and J_{RPC} on (b) CIFAR-10 and (c) LibriSpeech-100h with different minibatch sizes, showing that the performance of J_{RPC} is less sensitive to minibatch size change compared to J_{CPC} .

3.3 MINIBATCH SIZE SENSITIVITY

We then provide the analysis on the effect of minibatch size on J_{RPC} and J_{CPC} , since J_{CPC} is known to be sensitive to minibatch size (Poole et al., 2019). We train SimCLRv2 (Chen et al., 2020c) on CIFAR-10 and the model from Rife et al. (2020) on LibriSpeech-100h using J_{RPC} and J_{CPC} with different minibatch sizes. The settings of relative parameters are the same as Section 3.2. From Figure 1b and 1c, we can observe that both J_{RPC} and J_{CPC} achieve their optimal performance at a large minibatch size. However, when the minibatch size decreases, the performance of J_{RPC} shows higher sensitivity and suffers more when the number of minibatch samples is small. The result suggests that the proposed method might be less sensitive to the change of minibatch size compared to J_{CPC} given the same training settings.

3.4 EFFECT OF RELATIVE PARAMETERS

We study the effect of different combinations of relative parameters β , γ and α by comparing downstream performances on visual object recognition. We train SimCLRv2 on CIFAR-10 with different combinations of β , γ and α in J_{RPC} and α all other experimental settings. We choose $\beta \in \{0, 0.001, 1, 0\}$; $\gamma \in \{0, 0.001, 1, 0\}$; $\alpha \in \{0, 0.001, 1, 0\}$ and we report the best performances under each combination of β , γ and α . From Figure 2, we first observe that $\beta > 0$ has better downstream performance than $\beta = 0$ when γ and α are fixed. This observation is as expected, since $\beta > 0$ encourages representations of related and unrelated samples to be pushed away. Then, we find that a small but non-zero γ ($\gamma = 0.001$) and a large α ($\alpha = 1.0$) give the best performance compared to other combinations. Since β and γ serve as the coefficients of L_2 regularization, the results imply that the regularization is a strong and sensitive factor that will influence the performance. The results here are not as competitive as Table 2 because the CIFAR-10 result reported in Table 2 is using a set of relative parameters $(\beta = 1.0; \gamma = 0.005; \alpha = 1.0)$ that is different from the combinations in this subsection. Also, we use quite different ranges of imageNet (see A.7 for details.) In conclusion, we find empirically that a non-zero β , a small γ and a large α will lead to the optimal representation for the downstream task on CIFAR-10.

3.5 RELATION TO MUTUAL INFORMATION ESTIMATION

The presented approach also closely relates to mutual information estimation. For random variables X and Y with joint distribution P_{XY} and product of marginal $P_X P_Y$, the mutual information is defined as $I(X; Y) = D_{KL}(P_{XY} \| P_X P_Y)$. Lemma 1 states that given optimal solution for J_{RPC} , we can get the density ratio $r(x; y) := \frac{p(x; y)}{p(x)p(y)}$ as $r(x; y) = \frac{1 + \beta}{1 - \beta} \frac{p(x; y)}{p(x)p(y)}$. We can empirically estimate $r(x; y)$ from the estimated $\hat{r}(x; y)$ via this transformation, and use $\hat{r}(x; y)$ to estimate mutual information (Tsai et al., 2020b). Specifically, $I(X; Y) = \frac{1}{n} \sum_{i=1}^n \log r(x_i; y_i)$ with $(x_i; y_i) \sim P_{X; Y}^n$, where $P_{X; Y}^n$ is the uniformly sampled empirical distribution $P_{X; Y}^n$.

Figure 2: Heatmaps of downstream task performance on CIFAR-10, using different α and β in the J_{RPC} . We conclude that a nonzero α (small $\alpha = 0.001$) and a large β ($\beta = 1.0$) are crucial for better performance.

Figure 3: Mutual information estimation performed on 20-d correlated Gaussian distribution, with the correlation increasing each 4K steps. J_{RPC} exhibits smaller variance than SMILE and DoE, and smaller bias than J_{CPC} .

We follow prior work (Poole et al., 2019; Song & Ermon, 2019; Tsai et al., 2020b) for the experiments. We consider X and Y as two 20-dimensional Gaussians with correlation α and our goal is to estimate the mutual information $I(X; Y)$. Then, we perform a cubic transformation $y = 7! y^3$. The first task is referred to as a Gaussian task and the second is referred to as a Cubic task, where both have the ground truth $I(X; Y) = 10 \log(1 - \alpha^2)$. The models are trained 20,000 steps with $(X; Y)$ starting at $\alpha = 0$ and increased by $\alpha = 0.0004$ per 4,000 steps. Our method is compared with baseline methods J_{CPC} (Oord et al., 2018), J_{NWJ} (Nguyen et al., 2010), J_{JS} (Nowozin et al., 2016), SMILE (Song & Ermon, 2019) and Difference of Entropies (DoE) (McAllester & Stratos, 2020). All approaches use the same network design, learning rate, optimizer and minibatch size for a fair comparison. First, we observe J_{CPC} (Oord et al., 2018) has the smallest variance, while it exhibits a large bias (the estimated mutual information from J_{CPC} has an upper bound $\log(\text{batch size})$). Second, J_{NWJ} (Nguyen et al., 2010) and J_{JS} (Poole et al., 2019) have large variances, especially in the Cubic task. Song & Ermon (2019) pointed out the limitations of J_{CPC} , J_{NWJ} , and J_{JS} , and developed the SMILE method, which clips the value of the estimated density function to reduce the variance of the estimators. DoE (McAllester & Stratos, 2020) is neither a lower bound nor an upper bound of mutual information, but can achieve accurate estimates when underlying mutual information is large. J_{RPC} exhibits comparable bias and lower variance compared to the SMILE method, and is more stable than the DoE method. We would like to highlight our method’s low-variance property, where we neither clip the values of the estimated density ratio nor impose an upper bound on our estimated mutual information.

4 RELATED WORK

As a subset of unsupervised representation learning, self-supervised representation learning (SSL) adopts self-defined signals as supervision and uses the learned representation for downstream tasks, such as object detection and image captioning (Liu et al., 2020). We categorize SSL work into two groups: when the signal is the input’s hidden property or the corresponding view of the input. For the first group, for example, Jigsaw puzzle (Noroozi & Favaro, 2016) shuffles the image patches and defines the SSL task for predicting the shuffled positions of the image patches. Other instances are Predicting Rotations (Gidaris et al., 2018) and Shuffle & Learn (Misra et al., 2016). For the second group, the SSL task aims at modeling the co-occurrence of multiple views of data, via the contrastive or the predictive learning objectives (Tsai et al., 2020a). The predictive objectives encourage reconstruction from one view of the data to the other, such as predicting the lower part of an image from

its upper part (ImageGPT by [Chen et al. \(2020a\)](#)). Comparing the contrastive with predictive learning approaches, [Tsai et al. \(2020a\)](#) points out that the former requires less computational resources for a good performance but suffers more from the over-fitting problem.

Theoretical analysis ([Arora et al., 2019](#); [Tsai et al., 2020a](#); [Tosh et al., 2020](#)) suggests the contrastively learned representations can lead to a good downstream performance. Beyond the theory, [Tian et al. \(2020\)](#) shows what matters more for the performance are 1) the choice of the contrastive learning objective; and 2) the creation of the positive and negative data pairs in the contrastive objective. Recent work ([Khosla et al., 2020](#)) extends the usage of contrastive learning from the self-supervised setting to the supervised setting. The supervised setting defines the positive pairs as the data from the same class in the contrastive objective, while the self-supervised setting defines the positive pairs as the data with different augmentations.

Our work also closely relates to the skewed divergence measurement between distributions ([Lee, 1999](#); [2001](#); [Nielsen, 2010](#); [Yamada et al., 2013](#)). Recall that the usage of the relative parameters plays a crucial role to regularize our objective for its boundness and low variance. This idea is similar to the skewed divergence measurement, that when calculating the divergence between distributions P and Q , instead of considering $D(P \| Q)$, these approaches consider $D(\alpha P \| \alpha P + (1 - \alpha)Q)$ with D representing the divergence and $\alpha < 1$. A natural example is that the Jensen-Shannon divergence is a symmetric skewed KL divergence $D_{JS}(P \| Q) = 0.5D_{KL}(P \| 0.5P + 0.5Q) + 0.5D_{KL}(Q \| 0.5P + 0.5Q)$. Compared to the non-skewed counterpart, the skewed divergence has shown to have a more robust estimation for its value ([Lee, 1999](#); [2001](#); [Yamada et al., 2013](#)). Different from these works that focus on estimating the values of distribution divergence, we focus on learning self-supervised representations.

5 CONCLUSION

In this work, we present RPC, the Relative Predictive Coding, that achieves a good balance among the three challenges when modeling a contrastive learning objective: training stability, sensitivity to minibatch size, and downstream task performance. We believe this work brings an appealing option for training self-supervised models and inspires future work to design objectives for balancing the aforementioned three challenges. In the future, we are interested in applying RPC in other application domains and developing more principled approaches for better representation learning.

ACKNOWLEDGEMENT

This work was supported in part by the NSF IIS1763562, NSF Awards #1750439 #1722822, National Institutes of Health, IARPA D17PC00340, ONR Grant N000141812861, and Facebook PhD Fellowship. We would also like to acknowledge NVIDIA's GPU support and Cloud TPU support from Google's TensorFlow Research Cloud (TFRC).

REFERENCES

- Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensor flow: A system for large-scale machine learning. In 2016 USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), pp. 265–283, 2016.
- Martin Anthony and Peter L Bartlett. Neural network learning: Theoretical foundations. Cambridge university press, 2009.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. preprint arXiv:1902.09229, 2019.
- Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. arXiv preprint arXiv:2006.11477, 2020.

- Peter L Bartlett. The sample complexity of pattern classification with neural networks: the size of the weights is more important than the size of the network. *IEEE transactions on Information Theory* 44(2):525–536, 1998.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062* 2018.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882* 2020.
- Mark Chen, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, and Ilya Sutskever. Generative pretraining from pixels. *Proceedings of the 37th International Conference on Machine Learning* 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05702* 2020b.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029* 2020c.
- Ching-Yao Chuang, Joshua Robinson, Lin Yen-Chen, Antonio Torralba, and Stefanie Jegelka. De-biased contrastive learning. *arXiv preprint arXiv:2007.00224* 2020.
- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223, 2011.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* 2018.
- Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain markov process expectations for large time, *Communications on Pure and Applied Mathematics* 28(1):1–47, 1975.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728* 2018.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems* 30:5767–5777, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670* 2018.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation* 9(8):1735–1780, 1997.
- K Hornik, M Stinchcombe, and H White. Multilayer feedforward networks are universal approximators. *Neural Networks* 2(5):359–366, 1989.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* 2015.

- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. arXiv preprint arXiv:2004.11362, 2020.
- Thomas Kipf, Elise van der Pol, and Max Welling. Contrastive learning of structured world models. arXiv preprint arXiv:1911.12247, 2019.
- Lingpeng Kong, Cyprien de Masson d'Autume, Wang Ling, Lei Yu, Zihang Dai, and Dani Yogatama. A mutual information maximization perspective of language representation learning. arXiv preprint arXiv:1910.08350, 2019.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Lillian Lee. Measures of distributional similarity. In Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics, pp. 25–32, College Park, Maryland, USA, June 1999. Association for Computational Linguistics. doi: 10.3115/1034678.1034693. <https://www.aclweb.org/anthology/P99-1004>.
- Lillian Lee. On the effectiveness of the skew divergence for statistical language analysis. In TATS Citeseer, 2001.
- Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 510–519, 2019.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Zhaoyu Wang, Li Mian, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive? arXiv e-prints, pp. arXiv–2006, 2020.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In Proceedings of the IEEE international conference on computer vision, pp. 3730–3738, 2015.
- Sindy Löwe, Peter O'Connor, and Bastiaan Veeling. Putting an end to end-to-end: Gradient-isolated learning of representations. *Advances in Neural Information Processing Systems*, pp. 3039–3051, 2019.
- David McAllester and Karl Stratos. Formal limitations on the measurement of mutual information. In International Conference on Artificial Intelligence and Statistics, pp. 875–884, 2020.
- Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In European Conference on Computer Vision, pp. 527–544. Springer, 2016.
- XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- Frank Nielsen. A family of statistical symmetric divergences based on Jensen's inequality. arXiv preprint arXiv:1009.4004, 2010.
- Frank Nielsen and Richard Nock. On the chi square and higher-order chi distances for approximating f-divergences. *IEEE Signal Processing Letters*, 21(1):10–13, 2013.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In European Conference on Computer Vision, pp. 69–84. Springer, 2016.
- Sebastian Nowozin, Botond Cséke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, pp. 271–279, 2016.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748, 2018.

- Sherjil Ozair, Corey Lynch, Yoshua Bengio, Aaron Van den Oord, Sergey Levine, and Pierre Sermanet. Wasserstein dependency measure for representation learning. *Advances in Neural Information Processing Systems*, pp. 15604–15614, 2019.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 5206–5210. IEEE, 2015.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems* 8024–8035, 2019.
- Ben Poole, Sherjil Ozair, Aaron van den Oord, Alexander A Alemi, and George Tucker. On variational bounds of mutual information. *arXiv preprint arXiv:1905.06922*, 2019.
- Morgane Rivère, Armand Joulin, Pierre-Emmanuel Mazareau, and Emmanuel Dupoux. Unsupervised pretraining transfers well across languages. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 7414–7418. IEEE, 2020.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115(3):211–252, 2015.
- Jiaming Song and Stefano Ermon. Understanding the limitations of variational mutual information estimators. *arXiv preprint arXiv:1910.06222*, 2019.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. *arXiv preprint arXiv:2005.10243*, 2020.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. *arXiv preprint arXiv:2008.10150*, 2020.
- Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Demystifying self-supervised learning: An information-theoretical framework. *arXiv preprint arXiv:2006.05576*, 2020a.
- Yao-Hung Hubert Tsai, Han Zhao, Makoto Yamada, Louis-Philippe Morency, and Ruslan Salakhutdinov. Neural methods for point-wise dependency estimation. *arXiv preprint arXiv:2006.05553*, 2020b.
- Michael Tschannen, Josip Djolonga, Paul K Rubenstein, Sylvain Gelly, and Mario Lucic. On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*, 2019.
- Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Petar Velickovic, William Fedus, William L Hamilton, Pietro Li, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. *ICLR (Poster)* 2019.
- Makoto Yamada, Taiji Suzuki, Takafumi Kanamori, Hirotaka Hachiya, and Masashi Sugiyama. Relative density-ratio estimation for robust distribution comparison. *Neural computation* 25(5): 1324–1370, 2013.
- Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.

A APPENDIX

A.1 PROOF OF LEMMA 1 IN THE MAIN TEXT

Lemma 2 (Optimal Solution for J_{RPC} , restating Lemma 1 in the main text) Let

$$J_{RPC}(X; Y) := \sup_{f \in \mathcal{F}} \mathbb{E}_{P_{XY}} [f(x; y)] - \mathbb{E}_{P_X P_Y} [f(x; y)] - \frac{1}{2} \mathbb{E}_{P_{XY}} [f^2(x; y)] - \frac{1}{2} \mathbb{E}_{P_X P_Y} [f^2(x; y)]$$

and $r(x; y) = \frac{p(x; y)}{p(x)p(y)}$ be the density ratio. J_{RPC} has the optimal solution

$$f^*(x; y) = \frac{r(x; y)}{r(x; y) + 1} := r_{\leq 1/2}(x; y) \text{ with } -r_{\leq 1/2}(x; y) \leq \frac{1}{2}$$

Proof: The second-order functional derivative of the objective is

$$d^2 P_{X;Y} - d^2 P_X P_Y;$$

which is always negative. The negative second-order functional derivative implies the objective has a supreme value. Then, take the first-order functional derivative and set it to zero:

$$d P_{X;Y} - d P_X P_Y - f(x; y) d P_{X;Y} - f(x; y) d P_X P_Y = 0:$$

We then get

$$f^*(x; y) = \frac{d P_{X;Y}}{d P_{X;Y} + d P_X P_Y} = \frac{p(x; y) - p(x)p(y)}{p(x; y) + p(x)p(y)} = \frac{r(x; y) - 1}{r(x; y) + 1}$$

Since $0 \leq r(x; y) \leq 1$, we have $-\frac{1}{2} \leq \frac{r(x; y) - 1}{r(x; y) + 1} \leq \frac{1}{2}$. Hence,

$$f^*(x; y) \in [-\frac{1}{2}, \frac{1}{2}]; f^*(x; y) := r_{\leq 1/2}(x; y) \text{ with } -r_{\leq 1/2}(x; y) \leq \frac{1}{2}$$

A.2 RELATION BETWEEN J_{RPC} AND D_2

In this subsection, we aim to show the following: 1) $D_2(P_{XY} \parallel P_X P_Y) = \mathbb{E}_{P_X P_Y} [r^2(x; y)] - 1$; and 2) $J_{RPC}(X; Y) = -\frac{1}{2} \mathbb{E}_{P^0} [r^2(x; y)]$ by having $P^0 = \frac{1}{2} P_{XY} + \frac{1}{2} P_X P_Y$ as the mixture distribution of P_{XY} and $P_X P_Y$.

Lemma 3 $D_2(P_{XY} \parallel P_X P_Y) = \mathbb{E}_{P_X P_Y} [r^2(x; y)] - 1$

Proof: By definition (Nielsen & Nock, 2013),

$$\begin{aligned} D_2(P_{XY} \parallel P_X P_Y) &= \frac{\int \frac{d P_{XY}}{d P_X P_Y}^2}{\int \frac{d P_{XY}}{d P_X P_Y}} - 1 = \frac{\int \frac{d P_{XY}}{d P_X P_Y}^2}{\int \frac{d P_{XY}}{d P_X P_Y}} - 1 \\ &= \frac{\int \frac{p(x; y)}{p(x)p(y)}^2}{\int \frac{p(x; y)}{p(x)p(y)}} - 1 = \int r^2(x; y) d P_X P_Y - 1 \\ &= \mathbb{E}_{P_X P_Y} [r^2(x; y)] - 1 \end{aligned}$$

Lemma 4 Defining $P^0 = \frac{1}{2} P_{XY} + \frac{1}{2} P_X P_Y$ as a mixture distribution of P_{XY} and $P_X P_Y$, $J_{RPC}(X; Y) = -\frac{1}{2} \mathbb{E}_{P^0} [r^2(x; y)]$.

Proof: Plug in the optimal solution $\hat{h}(x; y) = \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y}$ (see Lemma 2) into J_{RPC} :

$$\begin{aligned} J_{RPC} &= \mathbb{E}_{P_{X;Y}} [f(x; y)] - \mathbb{E}_{P_X P_Y} [f(x; y)] - \frac{1}{2} \mathbb{E}_{P_{X;Y}} [f^2(x; y)] - \frac{1}{2} \mathbb{E}_{P_X P_Y} [f^2(x; y)] \\ &= \int_{\mathcal{Z}} f(x; y) \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} - \int_{\mathcal{Z}} f(x; y) dP_X P_Y - \frac{1}{2} \int_{\mathcal{Z}} f^2(x; y) \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} - \frac{1}{2} \int_{\mathcal{Z}} f^2(x; y) dP_X P_Y \\ &= \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} \frac{dP_X P_Y}{dP_X P_Y} - \int_{\mathcal{Z}} f(x; y) dP_X P_Y - \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} \frac{dP_X P_Y}{dP_X P_Y} - \frac{1}{2} \int_{\mathcal{Z}} f^2(x; y) dP_X P_Y \\ &= \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} - \int_{\mathcal{Z}} f(x; y) dP_X P_Y - \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} - \frac{1}{2} \int_{\mathcal{Z}} f^2(x; y) dP_X P_Y \end{aligned}$$

Since we define $r_{;;}(x; y) = \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y}$ and $P^0 = \frac{1}{2} P_{X;Y} + \frac{1}{2} P_X P_Y$,

$$J_{RPC} = \frac{1}{2} \mathbb{E}_{P^0} [r_{;;}^2(x; y)]:$$

A.3 PROOF OF PROPOSITION 1 IN THE MAIN TEXT

The proof contains two parts: showing $J_{RPC} \leq \frac{1}{2} + \frac{\epsilon^2}{2}$ (see Section A.3.1) and $J_{RPC}^{m;n}$ is a consistent estimator for J_{RPC} (see Section A.3.2).

A.3.1 BOUNDEDNESS OF J_{RPC}

Lemma 5 (Boundedness of J_{RPC}) $0 \leq J_{RPC} \leq \frac{1}{2} + \frac{\epsilon^2}{2}$

Proof: Lemma 4 suggests $J_{RPC}(X; Y) = \frac{1}{2} \mathbb{E}_{P^0} [r_{;;}^2(x; y)]$ with $P^0 = \frac{1}{2} P_{X;Y} + \frac{1}{2} P_X P_Y$ as the mixture distribution of $P_{X;Y}$ and $P_X P_Y$. Hence, it is obvious $J_{RPC}(X; Y) \geq 0$.

We leverage the intermediate results in the proof of Lemma 4:

$$\begin{aligned} J_{RPC}(X; Y) &= \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} \frac{dP_X P_Y}{dP_X P_Y} - \int_{\mathcal{Z}} f(x; y) dP_X P_Y - \frac{1}{2} \int_{\mathcal{Z}} \frac{dP_{X;Y}}{dP_{X;Y} + dP_X P_Y} - \frac{1}{2} \int_{\mathcal{Z}} f^2(x; y) dP_X P_Y \\ &= \frac{1}{2} \mathbb{E}_{P_{X;Y}} [r_{;;}(x; y)] - \mathbb{E}_{P_X P_Y} [f(x; y)] - \frac{1}{2} \mathbb{E}_{P_{X;Y}} [r_{;;}(x; y)] - \frac{1}{2} \mathbb{E}_{P_X P_Y} [f^2(x; y)] \end{aligned}$$

Since $r_{;;}(x; y) \in [0, 1]$, $J_{RPC}(X; Y) \leq \frac{1}{2} + \frac{\epsilon^2}{2}$.

A.3.2 CONSISTENCY

We first recall the definition of the estimation of J_{RPC} :

Definition 2 ($J_{RPC}^{m;n}$, empirical estimation of J_{RPC} , restating Definition 1 in the main text) We parameterize f via a family of neural networks $\mathcal{F} := \{f: \mathcal{Z} \rightarrow \mathbb{R}^d\}$ where $\mathcal{Z} \subseteq \mathbb{R}^d$ and \mathcal{F} is compact. Let $x_i, y_i, g_{j=1}^n$ be n samples drawn uniformly at random from $P_{X;Y}$ and $x_j^0, y_j^0, g_{j=1}^m$ be m samples drawn uniformly at random from $P_X P_Y$. Then,

$$J_{RPC}^{m;n} = \sup_{f \in \mathcal{F}} \left[\frac{1}{n} \sum_{i=1}^n f(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) - \frac{1}{n} \sum_{i=1}^n \frac{1}{2} f^2(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m \frac{1}{2} f^2(x_j^0; y_j^0) \right]$$

Our goal is to show that $\hat{J}_{RPC}^{m;n}$ is a consistent estimator for J_{RPC} . We begin with the following definition:

$$\hat{J}_{RPC}^{m;n} := \frac{1}{n} \sum_{i=1}^n f(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) - \frac{1}{n} \sum_{i=1}^n \frac{1}{2} f^2(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m \frac{1}{2} f^2(x_j^0; y_j^0) \quad (3)$$

and

$$E \hat{J}_{RPC}^{m;n} := E_{P_{XY}} [f(x; y)] - E_{P_X P_Y} [f(x; y)] - \frac{1}{2} E_{P_{XY}} [f^2(x; y)] - \frac{1}{2} E_{P_X P_Y} [f^2(x; y)] \quad (4)$$

Then, we follow the steps:

- The first part is about estimation. We show that, with high probability, $\hat{J}_{RPC}^{m;n}$ is close to $E \hat{J}_{RPC}^{m;n}$, for any given ϵ .
- The second part is about approximation. We will apply the universal approximation lemma of neural networks (Hornik et al., 1989) to show that there exists a network such that $E \hat{J}_{RPC}^{m;n}$ is close to J_{RPC} .

Part I - Estimation: With high probability, $\hat{J}_{RPC}^{m;n}$ is close to $E \hat{J}_{RPC}^{m;n}$, for any given ϵ .

Throughout the analysis on the uniform convergence, we need the assumptions on the boundness and smoothness of the function. Since we show the optimal function is bounded in J_{RPC} , we can use the same bounded values without losing too much precision. The smoothness of the function suggests that the output of the network should only change slightly when only slightly perturbing the parameters. Specifically, the two assumptions are as follows:

Assumption 1 (boundness of f) There exist universal constants such that $2F, C_U, C_L, f$ be the range of f and $U = \max_{j \in \{1, \dots, m\}} |C_U|, |C_L|, |f|$ be the maximal absolute value of f . In the paper, we can choose to constrain $C_U = 1$ and $C_L = -1$ since the optimal function has $-1 \leq f \leq 1$.

Assumption 2 (smoothness of f) There exists constant $\beta > 0$ such that $|f(x; y) - f(x'; y')| \leq \beta (\|x - x'\| + \|y - y'\|)$.

Now, we can bound the rate of uniform convergence of a function class in terms of covering number (Bartlett, 1998):

Lemma 6 (Estimation) Let $\epsilon > 0$ and $N(\cdot; \epsilon)$ be the covering number of with radius ϵ . Then,

$$\Pr \left(\sup_{f \in \mathcal{F}} |\hat{J}_{RPC}^{m;n} - E \hat{J}_{RPC}^{m;n}| \geq \epsilon \right) \leq 2N \left(\epsilon; \frac{1}{4(1 + \beta + 2(\beta + 1)U)} \right) \exp \left(-\frac{n \epsilon^2}{32M^2} \right) + \exp \left(-\frac{m \epsilon^2}{32M^2} \right) + \exp \left(-\frac{n \epsilon^2}{32U^2} \right) + \exp \left(-\frac{m \epsilon^2}{32U^2} \right)$$

Proof: For notation simplicity, we define the operators

- $P(f) = E_{P_{XY}} [f(x; y)]$ and $P_n(f) = \frac{1}{n} \sum_{i=1}^n f(x_i; y_i)$
- $Q(f) = E_{P_X P_Y} [f(x; y)]$ and $Q_m(f) = \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0)$

Hence,

$$\begin{aligned} \hat{J}_{RPC}^{m;n} - E \hat{J}_{RPC}^{m;n} &= P_n(f) - P(f) - Q_m(f) + Q(f) - P_n(f^2) + P(f^2) - Q_m(f^2) + Q(f^2) \\ &= (P_n(f) - P(f)) + (Q(f) - Q_m(f)) + (P_n(f^2) - P(f^2)) + (Q_m(f^2) - Q(f^2)) \end{aligned}$$

Let $\delta = \frac{\epsilon}{4 + 1 + 2(\epsilon + \epsilon)U}$ and $\mathcal{T} := \mathcal{N}(\epsilon; \mathcal{Q})$. Let $\mathcal{C} = \{f_1; f_2; \dots; f_T\}$ with $f_1; f_2; \dots; f_T \in \mathcal{G}$ be such that $B_1(\epsilon; \mathcal{Q})$, $B_1(\epsilon; \mathcal{Q})$ are δ -cover. Hence, for any $f \in \mathcal{F}$, there is $f_k \in \mathcal{C}$ such that $\|f - f_k\|_{\mathcal{Q}} \leq \delta$.

Then, for any $f_k \in \mathcal{C}$:

$$\begin{aligned} & \mathbb{E} \sup_{\text{RPC}}^{m,n} \mathcal{J}_{\text{RPC}}^h = \mathbb{E} \mathcal{J}_{\text{RPC}}^h \\ & \mathbb{E} |P_n(f) - P(f)| + \mathbb{E} |Q_m(f) - Q(f)| + |P_n(f^2) - P(f^2)| + |Q_m(f^2) - Q(f^2)| \\ & \mathbb{E} |P_n(f_k) - P(f_k)| + \mathbb{E} |P_n(f) - P_n(f_k)| + \mathbb{E} |P(f) - P(f_k)| \\ & + \mathbb{E} |Q_m(f_k) - Q(f_k)| + \mathbb{E} |Q_m(f) - Q_m(f_k)| + \mathbb{E} |Q(f) - Q(f_k)| \\ & + |P_n(f_k^2) - P(f_k^2)| + |P_n(f^2) - P_n(f_k^2)| + |P(f^2) - P(f_k^2)| \\ & + |Q_m(f_k^2) - Q(f_k^2)| + |Q_m(f^2) - Q_m(f_k^2)| + |Q(f^2) - Q(f_k^2)| \\ & \mathbb{E} |P_n(f_k) - P(f_k)| + \mathbb{E} |P_n(f) - P(f_k)| + \mathbb{E} |P(f) - P(f_k)| \\ & + \mathbb{E} |Q_m(f_k) - Q(f_k)| + \mathbb{E} |Q_m(f) - Q_m(f_k)| + \mathbb{E} |Q(f) - Q(f_k)| \\ & + |P_n(f_k^2) - P(f_k^2)| + 2U\epsilon + 2U\epsilon \\ & + |Q_m(f_k^2) - Q(f_k^2)| + 2U\epsilon + 2U\epsilon \\ & = \mathbb{E} |P_n(f_k) - P(f_k)| + \mathbb{E} |Q_m(f_k) - Q(f_k)| + |P_n(f_k^2) - P(f_k^2)| + |Q_m(f_k^2) - Q(f_k^2)| \\ & + 2\epsilon + 2\epsilon \\ & \mathbb{E} |P_n(f_k) - P(f_k)| + \mathbb{E} |Q_m(f_k) - Q(f_k)| + |P_n(f_k^2) - P(f_k^2)| + |Q_m(f_k^2) - Q(f_k^2)| + \frac{\epsilon}{2}; \end{aligned}$$

where

- $\mathbb{E} |P_n(f) - P_n(f_k)|$, $\mathbb{E} |P(f) - P(f_k)|$, $\mathbb{E} |Q_m(f) - Q_m(f_k)|$, and $\mathbb{E} |Q(f) - Q(f_k)|$.
- $|P_n(f^2) - P_n(f_k^2)| \leq 2kf_k$, $|P(f^2) - P(f_k^2)| \leq 2U\epsilon$ due to Assumptions 1 and 2. The result also applies for $|Q_m(f^2) - Q_m(f_k^2)|$, and $|Q(f^2) - Q(f_k^2)|$.

Hence,

$$\begin{aligned} & \Pr \left(\sup_{f \in \mathcal{F}} \mathcal{J}_{\text{RPC}}^{m,n} \geq \frac{\epsilon}{8} \right) \\ & \Pr \left(\max_{f_k \in \mathcal{C}} |P_n(f_k) - P(f_k)| + |Q_m(f_k) - Q(f_k)| + |P_n(f_k^2) - P(f_k^2)| + |Q_m(f_k^2) - Q(f_k^2)| + \frac{\epsilon}{2} \right) \\ & = \Pr \left(\max_{f_k \in \mathcal{C}} |P_n(f_k) - P(f_k)| + |Q_m(f_k) - Q(f_k)| + |P_n(f_k^2) - P(f_k^2)| + |Q_m(f_k^2) - Q(f_k^2)| \geq \frac{\epsilon}{2} \right) \\ & \stackrel{\text{X}^1}{=} \Pr \left(|P_n(f_k) - P(f_k)| + |Q_m(f_k) - Q(f_k)| + |P_n(f_k^2) - P(f_k^2)| + |Q_m(f_k^2) - Q(f_k^2)| \geq \frac{\epsilon}{2} \right) \\ & \stackrel{\text{X}^2}{=} \Pr \left(|P_n(f_k) - P(f_k)| \geq \frac{\epsilon}{8} \right) + \Pr \left(|Q_m(f_k) - Q(f_k)| \geq \frac{\epsilon}{8} \right) \\ & + \Pr \left(|P_n(f_k^2) - P(f_k^2)| \geq \frac{\epsilon}{8} \right) + \Pr \left(|Q_m(f_k^2) - Q(f_k^2)| \geq \frac{\epsilon}{8} \right); \end{aligned}$$

With Hoeffding's inequality,

- $\Pr \left(\left| P_n(f_k) - P(f_k) \right| \geq \frac{\epsilon}{8} \right) \leq 2 \exp \left(-\frac{n \epsilon^2}{32M^2} \right)$
- $\Pr \left(\left| Q_m(f_k) - Q(f_k) \right| \geq \frac{\epsilon}{8} \right) \leq 2 \exp \left(-\frac{m \epsilon^2}{32M^2} \right)$
- $\Pr \left(\left| P_n(f_k^2) - P(f_k^2) \right| \geq \frac{\epsilon}{8} \right) \leq 2 \exp \left(-\frac{n \epsilon^2}{32U^2} \right)$
- $\Pr \left(\left| Q_m(f_k^2) - Q(f_k^2) \right| \geq \frac{\epsilon}{8} \right) \leq 2 \exp \left(-\frac{m \epsilon^2}{32U^2} \right)$

To conclude,

$$\Pr \left(\sup_{f \in \mathcal{F}} \left| J_{RPC}^{m,n} - E J_{RPC} \right| \geq \frac{\epsilon}{2} \right) \leq 2N \left(\frac{\epsilon}{4}, \frac{1}{1 + 2(\frac{\epsilon}{4})^2} \right) \exp \left(-\frac{n \epsilon^2}{32M^2} \right) + \exp \left(-\frac{m \epsilon^2}{32M^2} \right) + \exp \left(-\frac{n \epsilon^2}{32U^2} \right) + \exp \left(-\frac{m \epsilon^2}{32U^2} \right)$$

Part II - Approximation: Neural Network Universal Approximation. We leverage the universal function approximation lemma of neural network

Lemma 7 (Approximation (Hornik et al., 1989)) Let $\epsilon > 0$. There exists $2N$ and a family of neural networks $\mathcal{F} := \{f : \mathbb{R}^d \rightarrow \mathbb{R}\}$ where \mathbb{R}^d is compact, such that $\inf_{f \in \mathcal{F}} E J_{RPC} = J_{RPC} + \epsilon$.

Part III - Bringing everything together. Now, we are ready to bring the estimation and approximation together to show that there exists a neural network such that, with high probability, can approximate J_{RPC} with $n^0 = \min\{n, m\}$ at a rate of $O(1/n^0)$:

Proposition 3 With probability at least $1 - \frac{\epsilon}{2}$, $\left| J_{RPC} - J_{RPC}^{m,n} \right| \leq O \left(\frac{1}{n^0} \right)$; where $n^0 = \min\{n, m\}$.

Proof: The proof follows by combining Lemma 6 and 7.

First, Lemma 7 suggests,

$$E J_{RPC} - J_{RPC} \leq \frac{\epsilon}{2}$$

Next, we perform analysis on the estimation error, aiming to find the corresponding probability, such that

$$J_{RPC}^{m,n} - E J_{RPC} \leq \frac{\epsilon}{2}$$

Applying Lemma 6 with the covering number of the neural network: $N(\cdot, \cdot) =$

$O \left(\exp(d \log(1/\epsilon)) \right)$ (Anthony & Bartlett, 2009) and let $n^0 = \min\{n, m\}$:

$$\Pr \left(\sup_{f \in \mathcal{F}} \left| J_{RPC}^{m,n} - E J_{RPC} \right| \geq \frac{\epsilon}{2} \right) \leq 2N \left(\frac{\epsilon}{8}, \frac{1}{1 + 2(\frac{\epsilon}{8})^2} \right) \exp \left(-\frac{n \epsilon^2}{128M^2} \right) + \exp \left(-\frac{m \epsilon^2}{128M^2} \right) + \exp \left(-\frac{n \epsilon^2}{128U^2} \right) + \exp \left(-\frac{m \epsilon^2}{128U^2} \right)$$

$$= O \left(\exp(d \log(1/\epsilon)) \frac{1}{n^0} \right)$$

where the big-O notation absorbs all the constants that do not require in the following derivation. Since we want to bound the probability with ϵ , we solve the such that

$$\exp(-d \log(1-\epsilon)) \leq n^{-0.2} \epsilon.$$

With $\log(x) \leq -x + 1$,

$$n^{-0.2} + d(1-\epsilon) \leq n^{-0.2} + d \log(1-\epsilon);$$

where this inequality holds when

$$\epsilon = O\left(\frac{d + \log(1-\epsilon)}{n^{0.2}}\right).$$

A.4 PROOF OF PROPOSITION 2 IN THE MAIN TEXT - FROM AN ASYMPTOTIC VIEWPOINT

Here, we provide the variance analysis of $\hat{J}_{RPC}^{m;n}$ via an asymptotic viewpoint. First, assuming the network is correctly specified, and hence there exists a network parameter satisfying $f^*(x; y) = f^*(x; y) = r^*(x; y)$. Then we recall that $\hat{J}_{RPC}^{m;n}$ is a consistent estimator of $J_{RPC}^{m;n}$ (see Proposition 3), and under regular conditions, the estimated network parameter $\hat{\theta}_{RPC}^{m;n}$ satisfying the asymptotic normality in the large sample limit (see Theorem 5.23 in (Van der Vaart, 2000)). We recall the definition of $J_{RPC}^{m;n}$ in equation 3 and let $\epsilon_0 = \min_{f \in \mathcal{F}} \min_{g \in \mathcal{G}}$, the asymptotic expansion of $\hat{J}_{RPC}^{m;n}$ has

$$\begin{aligned} \hat{J}_{RPC}^{m;n} &= J_{RPC}^{m;n} + \hat{J}_{RPC}^{m;n}(\hat{\theta}) + o_p(k^{-k}) \\ &= J_{RPC}^{m;n} + \hat{J}_{RPC}^{m;n}(\hat{\theta}) + o_p\left(\frac{1}{n^0}\right) \\ &= J_{RPC}^{m;n} + o_p\left(\frac{1}{n^0}\right); \end{aligned} \tag{5}$$

where $\hat{J}_{RPC}^{m;n}(\hat{\theta}) = 0$ since $\hat{\theta}$ is the estimation from $\hat{J}_{RPC}^{m;n} = \sup_{f \in \mathcal{F}} J_{RPC}^{m;n}$.

Next, we recall the definition in equation 4:

$$E[\hat{J}_{RPC}^{m;n}(\hat{\theta})] = E_{P_{XY}}[f^*(x; y)] - E_{P_X P_Y}[f^*(x; y)] = \frac{1}{2} E_{P_{XY}}[f^*(x; y)] - \frac{1}{2} E_{P_X P_Y}[f^*(x; y)];$$

Likewise, the asymptotic expansion of $E[\hat{J}_{RPC}^{m;n}(\hat{\theta})]$ has

$$\begin{aligned} E[\hat{J}_{RPC}^{m;n}(\hat{\theta})] &= E[J_{RPC}^{m;n}] + E[\hat{J}_{RPC}^{m;n}(\hat{\theta})] + o_p(k^{-k}) \\ &= E[J_{RPC}^{m;n}] + E[\hat{J}_{RPC}^{m;n}(\hat{\theta})] + o_p\left(\frac{1}{n^0}\right) \\ &= E[J_{RPC}^{m;n}] + o_p\left(\frac{1}{n^0}\right); \end{aligned} \tag{6}$$

where $E[\hat{J}_{RPC}^{m;n}(\hat{\theta})] = 0$ since $E[J_{RPC}^{m;n}] = J_{RPC}$ and $\hat{\theta}$ satisfying $f^*(x; y) = f^*(x; y)$.

Combining equations 5 and 6:

$$\begin{aligned}
\mathbb{J}_{\text{RPC}; \hat{\Lambda}}^{m;n} \quad \mathbb{E}[\hat{\mathbb{J}}_{\text{RPC}; \hat{\Lambda}}^{m;n}] &= \mathbb{J}_{\text{RPC}}^{m;n} + \mathbb{J}_{\text{RPC}} + o_p\left(\frac{1}{n^0}\right) \\
&= \frac{1}{n} \sum_{i=1}^n f(x_i; y_i) + \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) - \frac{1}{2n} \sum_{i=1}^n f^2(x_i; y_i) - \frac{1}{2m} \sum_{j=1}^m f^2(x_j^0; y_j^0) \\
&\quad \mathbb{E}_{P_{XY}} [f(x; y)] + \mathbb{E}_{P_X P_Y} [f(x; y)] + \frac{1}{2} \mathbb{E}_{P_{XY}} [f^2(x; y)] + \frac{1}{2} \mathbb{E}_{P_X P_Y} [f^2(x; y)] + o_p\left(\frac{1}{n^0}\right) \\
&= \frac{1}{n} \sum_{i=1}^n r_{i; ;} (x_i; y_i) + \frac{1}{m} \sum_{j=1}^m r_{j; ;} (x_j^0; y_j^0) - \frac{1}{2n} \sum_{i=1}^n r_{i; ;}^2 (x_i; y_i) - \frac{1}{2m} \sum_{j=1}^m r_{j; ;}^2 (x_j^0; y_j^0) \\
&\quad \mathbb{E}_{P_{XY}} [r_{i; ;} (x; y)] + \mathbb{E}_{P_X P_Y} [r_{j; ;} (x; y)] + \frac{1}{2} \mathbb{E}_{P_{XY}} [r_{i; ;}^2 (x; y)] + \frac{1}{2} \mathbb{E}_{P_X P_Y} [r_{j; ;}^2 (x; y)] \\
&\quad + o_p\left(\frac{1}{n^0}\right) \\
&= \frac{1}{n} \sum_{i=1}^n r_{i; ;} (x_i; y_i) - \frac{1}{2n} \sum_{i=1}^n r_{i; ;}^2 (x_i; y_i) + \mathbb{E}_{P_{XY}} [r_{i; ;} (x; y)] - \frac{1}{2} \mathbb{E}_{P_{XY}} [r_{i; ;}^2 (x; y)] \\
&\quad + \frac{1}{m} \sum_{j=1}^m r_{j; ;} (x_j^0; y_j^0) + \frac{1}{2m} \sum_{j=1}^m r_{j; ;}^2 (x_j^0; y_j^0) + \mathbb{E}_{P_X P_Y} [r_{j; ;} (x; y)] + \frac{1}{2} \mathbb{E}_{P_X P_Y} [r_{j; ;}^2 (x; y)] \\
&\quad + o_p\left(\frac{1}{n^0}\right):
\end{aligned}$$

Therefore, the asymptotic Variance of $\mathbb{J}_{\text{RPC}}^{m;n}$ is

$$\text{Var}[\mathbb{J}_{\text{RPC}}^{m;n}] = \frac{1}{n} \text{Var}_{P_{XY}} [r_{i; ;} (x; y) - \frac{1}{2} r_{i; ;}^2 (x; y)] + \frac{1}{m} \text{Var}_{P_X P_Y} [r_{j; ;} (x; y) + \frac{1}{2} r_{j; ;}^2 (x; y)] + o\left(\frac{1}{n^0}\right)$$

First, we look at $\text{Var}_{P_{XY}} [r_{i; ;} (x; y) - \frac{1}{2} r_{i; ;}^2 (x; y)]$. Since $r_{i; ;} > 0$ and $-\frac{1}{2} r_{i; ;}^2 < 0$, simple calculation gives us $\frac{2}{2^2} r_{i; ;}^2 (x; y) + \frac{1}{2} r_{i; ;}^2 (x; y) = \frac{1}{2} r_{i; ;}^2 (x; y)$. Hence,

$$\text{Var}_{P_{XY}} [r_{i; ;} (x; y) - \frac{1}{2} r_{i; ;}^2 (x; y)] = \max \left\{ \frac{2}{2^2} r_{i; ;}^2 (x; y); \frac{1}{2} r_{i; ;}^2 (x; y) \right\}$$

Next, we look at $\text{Var}_{P_X P_Y} [r_{j; ;} (x; y) + \frac{1}{2} r_{j; ;}^2 (x; y)]$. Since $r_{j; ;} > 0$ and $\frac{1}{2} r_{j; ;}^2 > 0$, simple calculation gives us $\frac{2}{2^2} r_{j; ;}^2 (x; y) + \frac{1}{2} r_{j; ;}^2 (x; y) = \frac{2}{2^2} r_{j; ;}^2 (x; y)$. Hence,

$$\text{Var}_{P_X P_Y} [r_{j; ;} (x; y) + \frac{1}{2} r_{j; ;}^2 (x; y)] = \max \left\{ \frac{2}{2^2} r_{j; ;}^2 (x; y); \frac{2}{2^2} r_{j; ;}^2 (x; y) \right\}$$

Combining everything together, we restate the Proposition 2 in the main text:

Proposition 4 (Asymptotic Variance of $\mathbb{J}_{\text{RPC}}^{m;n}$)

$$\begin{aligned}
\text{Var}[\mathbb{J}_{\text{RPC}}^{m;n}] &= \frac{1}{n} \text{Var}_{P_{XY}} [r_{i; ;} (x; y) - \frac{1}{2} r_{i; ;}^2 (x; y)] + \frac{1}{m} \text{Var}_{P_X P_Y} [r_{j; ;} (x; y) + \frac{1}{2} r_{j; ;}^2 (x; y)] + o\left(\frac{1}{n^0}\right) \\
&= \frac{1}{n} \max \left\{ \frac{2}{2^2} r_{i; ;}^2 (x; y); \frac{1}{2} r_{i; ;}^2 (x; y) \right\} + \frac{1}{m} \max \left\{ \frac{2}{2^2} r_{j; ;}^2 (x; y); \frac{2}{2^2} r_{j; ;}^2 (x; y) \right\} + o\left(\frac{1}{n^0}\right)
\end{aligned}$$

A.5 PROOF OF PROPOSITION 2 IN THE MAIN TEXT - FROM BOUNDEDNESS OF

As discussed in Assumption 1, for the estimation of $J_{RPC}^{m;n}$, we can bound the function in F within $[-1, 1]$ without losing precision. Then, re-arranging:

$$\begin{aligned} \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n f(x_i; y_i) &= \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) + \frac{1}{n} \sum_{i=1}^n \frac{1}{2} f^2(x_i; y_i) - \frac{1}{m} \sum_{j=1}^m \frac{1}{2} f^2(x_j^0; y_j^0) \\ \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n f(x_i; y_i) &= \frac{1}{2} \sum_{i=1}^n f^2(x_i; y_i) + \frac{1}{m} \sum_{j=1}^m f(x_j^0; y_j^0) + \frac{1}{2} \sum_{j=1}^m f^2(x_j^0; y_j^0) \end{aligned}$$

Then, since $f(x; y) \in [-1, 1]$, basic calculations give us

$$\frac{2 + \frac{1}{2}}{2} f(x_i; y_i) - \frac{1}{2} f^2(x_i; y_i) \leq \frac{1}{2} \text{ and } \frac{1}{2} f(x_j^0; y_j^0) + \frac{1}{2} f^2(x_j^0; y_j^0) \leq \frac{2 + \frac{1}{2}}{2} :$$

The resulting variances have

$$\text{Var}[f(x_i; y_i) - \frac{1}{2} f^2(x_i; y_i)] \leq \max\left\{\frac{2 + \frac{1}{2}}{2}, \frac{1}{2}\right\}^2$$

and

$$\text{Var}[f(x_j^0; y_j^0) + \frac{1}{2} f^2(x_j^0; y_j^0)] \leq \max\left\{\frac{2 + \frac{1}{2}}{2}, \frac{2 + \frac{1}{2}}{2}\right\}^2 :$$

Taking the mean of n independent random variables gives the result:

Proposition 5 (Variance of $J_{RPC}^{m;n}$)

$$\text{Var}[J_{RPC}^{m;n}] \leq \frac{1}{n} \max\left\{\frac{2 + \frac{1}{2}}{2}, \frac{1}{2}\right\}^2 + \frac{1}{m} \max\left\{\frac{2 + \frac{1}{2}}{2}, \frac{2 + \frac{1}{2}}{2}\right\}^2 :$$

A.6 IMPLEMENTATION OF EXPERIMENTS

For visual representation learning, we follow the implementation <https://github.com/google-research/simclr>. For speech representation learning, we follow the implementation in https://github.com/facebookresearch/CPC_audio. For MI estimation, we follow the implementation in https://github.com/yaohungt/Pointwise_Dependency_Neural_Estimation/tree/master/MI_Est_and_CrossModal.

A.7 RELATIVE PREDICTIVE CODING ON VISION

The whole pipeline of pretraining contains the following steps: First, a stochastic data augmentation will transform one image sample x_k to two different but correlated augmented views x_{2k-1}^0 and x_{2k}^0 . Then a base encoder $f(\cdot)$ implemented using ResNet (He et al., 2016) will extract representations from augmented views, creating representations z_{2k-1} and z_{2k} . Later a small neural network $g(\cdot)$ called projection head will map z_{2k-1} and z_{2k} to z_{2k-1} and z_{2k} in a different latent space. For each minibatch of N samples, there will be $2N$ views generated. For each image there will be one positive pair x_{2k-1}^0 and x_{2k}^0 and $2(N-1)$ negative samples. The RPC loss between a pair of positive views x_i^0 and x_j^0 (augmented from the same image), can be calculated by the substitution $f(x_i^0; x_j^0) = (z_i - z_j) = s_{ij}$ (s is a hyperparameter) to the definition of RPC:

$$J_{ij}^{RPC} = (s_{ij} - \frac{1}{2(N-1)} \sum_{k=1}^N \mathbb{1}_{[k \neq i]} s_{i;k} - \frac{1}{2} s_{ij}^2 - \frac{1}{2(N-1)} \sum_{k=1}^N \mathbb{1}_{[k \neq i]} s_{i;k}^2) \quad (7)$$

For losses other than RPC, a hidden normalization is often required by replacing $(z_i - z_j)$ with $(z_i - z_j) / \sqrt{z_i z_j}$. CPC and WPC adopt this, while other objectives need it to help stabilize training variance. RPC does not need this normalization.

Con dence Interval of J_{RPC} and J_{CPC}			
Objective	CIFAR 10	CIFAR 100	ImageNet
J_{CPC}	(91:09%; 91:13%)	(77:11%; 77:36%)	(73:39%; 73:48%)
J_{RPC}	(91:16%; 91:47%)	(77:41%; 77:98%)	(73:92%; 74:43%)

Table 4: Con dence Intervals of performances J_{RPC} and J_{CPC} on CIFAR-10/-100 and ImageNet.

A.8 CIFAR-10/-100 AND IMAGENET EXPERIMENTS DETAILS

ImageNet Following the settings in (Chen et al., 2020b;c), we train the model on Cloud TPU with 128 cores, with a batch size of 096 and global batch normalization (Ioffe & Szegedy, 2015). Here we refer to the term batch size as the number of images (or utterances in the speech experiments) we use per GPU, while the term minibatch size refers to the number of negative samples used to calculate the objective, such as CPC or our proposed RPC. The largest model we train is a 152-layer ResNet with selective kernels (SK) (Li et al., 2019) and wider channels. We use the LARS optimizer (You et al., 2017) with momentum 0.9. The learning rate linearly increases for the first 20 epochs, reaching a maximum of 0.04, then decayed with cosine decay schedule. The weight decay is 10^{-4} . A MLP projection head () with three layers is used on top of the ResNet encoder. Unlike Chen et al. (2020c), we do not use a memory buffer, and train the model for 1000 epochs rather than 800 epochs due to computational constraints. These two options slightly reduce CPC’s performance benchmark for about 2% with the exact same setting. The unsupervised pre-training is followed by a supervised re-tuning. Following SimCLRv2 (Chen et al., 2020b;c), we re-tune the 3-layer () for the downstream tasks. We use learning rates 0.04 and 0.064 for standard 50-layer ResNet and larger 152-layer ResNet respectively, and weight decay and learning rate warmup are removed. Different from Chen et al. (2020c), we use a batch size of 096, and we do not use global batch normalization for re-tuning. For J_{RPC} we disable hidden normalization and use a temperature = 32. For all other objectives, we use hidden normalization and 0.1 following previous work (Chen et al., 2020c). For relative parameters, we use $\alpha = 0.3$; $\beta = 0.001$; $\gamma = 0.1$ and $\delta = 0.3$; $\epsilon = 0.001$; $\zeta = 0.005$ for ResNet-50 and ResNet-152 respectively.

CIFAR-10/-100 Following the settings in (Chen et al., 2020b), we train the model on a single GPU, with a batch size of 12 and global batch normalization (Ioffe & Szegedy, 2015). We use ResNet (He et al., 2016) of depth 18 and depth 50, and does not use Selective Kernel (Li et al., 2019) or a multiplied width size. We use the LARS optimizer (You et al., 2017) with momentum 0.9. The learning rate linearly increases for the first 20 epochs, reaching a maximum of 0.04, then decayed with cosine decay schedule. The weight decay is 10^{-4} . A MLP projection head () with three layers is used on top of the ResNet encoder. Unlike Chen et al. (2020c), we do not use a memory buffer. We train the model for 1000 epochs. The unsupervised pre-training is followed by a supervised re-tuning. Following SimCLRv2 (Chen et al., 2020b;c), we re-tune the 3-layer () for the downstream tasks. We use learning rates 0.04 for standard 50-layer ResNet, and weight decay and learning rate warmup are removed. For J_{RPC} we disable hidden normalization and use a temperature = 128. For all other objectives, we use hidden normalization and 0.5 following previous work (Chen et al., 2020c). For relative parameters, we use $\alpha = 0$; $\beta = 0.005$; and $\gamma = 1.0$.

STL-10 We also perform the pre-training and re-tuning on STL-10 (Coates et al., 2011) using the model proposed in Chuang et al. (2020). Chuang et al. (2020) proposed to indirectly approximate the distribution of negative samples so that the objective is unbiased. However, their implementation of contrastive learning is consistent with Chen et al. (2020b). We use a ResNet with 50-layer encoder for pre-training, with Adam optimizer, learning rate 0.01 and weight decay 10^{-6} . The temperature is set to 0.5 for all objectives other than J_{RPC} , which disables hidden normalization and use = 128. The downstream task performance increases 83.4% of J_{CPC} to 84.1% of J_{RPC} .

Con dence Interval We also provide the con dence interval of J_{RPC} and J_{CPC} on CIFAR-10, CIFAR-100 and ImageNet, using ResNet-18, ResNet-18 and ResNet-50 respectively (95% con -

³For WPC (Ozair et al., 2019), the global batch normalization during pretraining is disabled since we enforce 1-Lipschitz by gradient penalty (Gulrajani et al., 2017).

dence level is chosen) in Table 4. Both CPC and RPC use the same experimental settings throughout this paper. Here we use the relative parameters $(\beta : 0; \gamma = 0.005; \tau = 1.0)$ in J_{RPC} which gives the best performance on CIFAR-10. The confidence intervals of CPC do not overlap with the confidence intervals of RPC, which means the difference of the downstream task performance between RPC and CPC is statistically significant.

A.9 RELATIVE PREDICTIVE CODING ON SPEECH

For speech representation learning, we adopt the general architecture from Oord et al. (2018). Given an input signal $x_{1:T}$ with T time steps, we first pass it through an encoder parametrized by ϕ to produce a sequence of hidden representations h_1, \dots, h_T where $h_t = \phi(x_t)$. After that, we obtain the contextual representation c_t at time step t with a sequential model parametrized by ψ : $c_t = \psi(h_1, \dots, h_t)$, where c_t contains context information before time step t . For unsupervised pre-training, we use a multi-layer convolutional network as the encoder and an LSTM with hidden dimension 256 as the sequential model. Here, the contrastiveness is between the positive pair $(h_{t+k}; c_t)$ where k is the number of time steps ahead, and the negative pairs $(h_i; c_t)$, where h_i is randomly sampled from \mathcal{H} , a batch of hidden representation of signals assumed to be unrelated to c_t . The scoring function based on Equation 2 at step and look-ahead k will be $f_k = f_k(h; c_t) = \exp(\langle h \rangle^T W_k c_t)$, where W_k is a learnable linear transformation defined separately for each f_1, \dots, f_K and K is predetermined at 2 time steps. The loss in Equation 2 will then be formulated as:

$$J_{t,k}^{RPC} = \frac{1}{|N_j|} \sum_{h_i \in 2N} f_k(h_{t+k}; c_t) - \frac{1}{2} \frac{f_k^2(h_{t+k}; c_t)}{|N_j|} - \frac{1}{2|N_j|} \sum_{h_i \in 2N} f_k^2(h_i; c_t) \quad (8)$$

We use the following relative parameters $\beta = 1$; $\gamma = 0.25$; and $\tau = 1$, and we use the temperature $\beta = 16$ for J_{RPC} . For J_{CPC} we follow the original implementation which sets $\beta = 1$. We fix all other experimental setups, including architecture, learning rate, and optimizer. As shown in Table 3, J_{RPC} has better downstream task performance, and is closer to the performance from a fully supervised model.

A.10 EMPIRICAL OBSERVATIONS ON VARIANCE AND MINIBATCH SIZE

Variance Experiment Setup We perform the variance comparison of J_{DV} , J_{NWJ} and the proposed J_{RPC} . The empirical experiments are performed using SimCLRv2 (Chen et al., 2020c) on CIFAR-10 dataset. We use a ResNet of depth 18 with batch size of 512. We train each objective with 30K training steps and record their value. In Figure 1, we use a temperature of 28 for all objectives. Unlike other experiments, where hidden normalization is applied to other objectives, we remove hidden normalization for all objectives due to the reality that objectives after normalization does not reflect their original values. From Figure 1, J_{RPC} enjoys lower variance and more stable training compared to J_{DV} and J_{NWJ} .

Minibatch Size Experimental Setup We perform experiments on the effect of batch size on downstream performances for different objective. The experiments are performed using SimCLRv2 (Chen et al., 2020c) on CIFAR-10 dataset, as well as the model from Rivest et al. (2020) on LibriSpeech-100h dataset (Panayotov et al., 2015). For vision task, we use the default temperature $\beta = 0.5$ from Chen et al. (2020c) and hidden normalization mentioned in Section 3. For J_{RPC} in vision and speech tasks we use a temperature of 28 and $\beta = 16$ respectively, both without hidden normalization.

A.11 MUTUAL INFORMATION ESTIMATION

Our method is compared with baseline methods CPC (Oord et al., 2018), NWJ (Nguyen et al., 2010), JSD (Nowozin et al., 2016), and SMILE (Song & Ermon, 2019). All the approaches consider the same design of $\phi(x; y)$, which is a 3-layer neural network taking concatenated $(x; y)$ as the input. We also fix the learning rate, the optimizer, and the minibatch size across all the estimators for a fair comparison.

We present results of mutual information by Relative Predictive Coding using different sets of relative parameters in Figure 4. In the first row, we set $\beta = 10^{-3}$, $\gamma = 1$, and experiment with different

