

Efficient Maximal Coding Rate Reduction by Variational Forms

Christina Baek*

Carnegie Mellon University

kbaek@cs.cmu.edu *equal contribution

Ziyang Wu*

International Digital Economy Academy

wuziyang@idea.edu.cn *equal contribution

Kwan Ho Ryan Chan

Johns Hopkins University

kchan49@jhu.edu

Tianjiao Ding

Johns Hopkins University

tding@jhu.edu

Yi Ma

University of California, Berkeley

yima@eecs.berkeley.edu

Benjamin D. Haeffele

Johns Hopkins University

bhaeffele@jhu.edu

Abstract

The principle of Maximal Coding Rate Reduction (MCR²) has recently been proposed as a training objective for learning discriminative low-dimensional structures intrinsic to high-dimensional data to allow for more robust training than standard approaches, such as cross-entropy minimization. However, despite the advantages that have been shown for MCR² training, MCR² suffers from a significant computational cost due to the need to evaluate and differentiate a significant number of log-determinant terms that grows linearly with the number of classes. By taking advantage of variational forms of spectral functions of a matrix, we reformulate the MCR² objective to a form that can scale significantly without compromising training accuracy. Experiments in image classification demonstrate that our proposed formulation results in a significant speed up over optimizing the original MCR² objective directly and often results in higher quality learned representations. Further, our approach may be of independent interest in other models that require computation of log-determinant forms, such as in system identification or normalizing flow models.

1. Introduction

Given a classification task, deep networks aim to learn a nonlinear mapping, consisting of a series of linear and nonlinear functions, that can map data to their correct labels. The overall deep network can often be interpreted as a composition of a nonlinear “featurizer” f_θ and a linear classifier $g(z) = \mathbf{W}z$ for some matrix \mathbf{W} . The hidden layers or

the *featurizer* is designated with learning a latent representation $\mathbf{z}_\theta = f_\theta(\mathbf{x}) \in \mathbb{R}^d$ that best facilitates the final layer or *classifier* for the downstream task.

The canonical way to train a deep learning model for a classification task is empirical risk minimization using cross-entropy (CE) loss. While CE measures the difference between the model’s prediction and the true labels, it does not explicitly enforce any structure over the representation. In fact, Papayan, Han, and Donoho [8, 17] show that this direct label fitting implicitly leads to *neural collapse* in deep networks. That is, as CE loss converges to 0, the representations of each class at the last hidden layer collapse to a single point, suppressing within-class variability.

Beyond neural collapse and failing to represent within-class variation, several works [1, 5, 19] have empirically shown that training neural networks using stochastic gradient descent (SGD) on CE loss often leads the network to utilize the simplest, often spurious, feature in the image for classification. This hypothesis is theoretically supported by [2] which verified that when multiple explanations can describe a class, models trained with CE often pick a subset of features that can classify a majority of the points well and then classifies the remaining points from noise in the data.

To alleviate this issue, Yu et al. [21] proposed a framework for learning geometrically meaningful representations, via a featurizer f_θ , by maximizing the coding rate reduction (MCR²). In brief, the MCR² objective encourages the latent representation of the entire training set to *expand* or occupy as much volume as possible, while simultaneously pushing each class to *compress* or occupy as little space as possible. Empirically and theoretically, it is shown that this objective drives the latent representations of

each class to lie on a low dimensional *linear* subspace, with the subspaces orthogonal to each other, which empirically provides robustness against label noise, a notable advantage of MCR² compared to CE [21].

However, despite these inherent advantages, the MCR² metric can be costly from a computational perspective. In particular, the loss involves calculating the log det of the Gram matrix of the representations of each class. Not only does the number of log det terms grow linearly with the number of classes, but computing (and back-propagating) the log det of a $d \times d$ matrix incurs a computational cost of $O(d^3)$. For this reason, MCR² methods to date have been limited to datasets with a relatively small number of classes such as MNIST and CIFAR-10, where the loss is computationally feasible. In order to make MCR² scalable, there is a significant need to improve the computational efficiency, particularly as it pertains to computing log det terms, to allow for large numbers (hundreds or thousands) of classes in high dimensional spaces.

Contributions. In this paper we make significant progress towards this goal, with the following contributions:

1. We provide an alternative formulation of the MCR² objective based on a variational form of the log det function which scales much more gracefully with the number of classes and the problem dimension.
2. We show experimentally that the alternative formulation requires approximately the same number of epochs to converge as the original MCR² formulation, but achieves a significant speedup in the training cost per epoch, particularly as the number of classes in the dataset grows.
3. As an additional benefit of our formulation, we also observe empirically that training over our proposed variational formulation often results in higher quality learned latent representations and better test accuracy than the original MCR² objective.

Finally, we note that our approach for optimization with variational forms may be of independent interest for other models which require computing log det terms, such as in system identification [4] or normalizing flow models [9].

2. Preliminaries

Here we first describe the original MCR² formulation as well as introduce relevant background material.

2.1. MCR² Objective

The original MCR² objective [21] takes the following form: Given m training samples $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_m] \in \mathbb{R}^{D \times m}$ belonging to k classes¹, let $\mathbf{Z}_\theta =$

¹Here we adopt the notation that an upper case letter \mathbf{X} represents a matrix and \mathbf{X}_i denotes the i^{th} column of a matrix.

$[f_\theta(\mathbf{X}_1), \dots, f_\theta(\mathbf{X}_m)] \in \mathbb{R}^{d \times m}$ be the latent representation where recall f_θ is the featurizer parameterized by θ , and let $\mathbf{\Pi} \in \mathbb{R}^{m \times k}$ define the class membership, where $\Pi_{i,j}$ denotes the probability² that \mathbf{X}_i is in class j . Then, MCR² aims to learn a feature representation \mathbf{Z}_θ that maximizes the following *coding rate reduction* $\Delta R(\mathbf{Z}_\theta)$:

$$\begin{aligned} \max_{\theta} \Delta R(\mathbf{Z}_\theta) &\equiv R(\mathbf{Z}_\theta) - R_c(\mathbf{Z}_\theta, \mathbf{\Pi}) \quad \text{s.t. } \mathbf{Z}_\theta \in \mathcal{S}, \quad \text{where} \\ R(\mathbf{Z}_\theta) &= \frac{1}{2} \log \det \left(\mathbf{I} + \alpha \mathbf{Z}_\theta \mathbf{Z}_\theta^\top \right), \quad \text{and} \\ R_c(\mathbf{Z}_\theta, \mathbf{\Pi}) &= \sum_{j=1}^k \frac{\gamma_j}{2} \log \det \left(\mathbf{I} + \alpha_j \mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top \right) \end{aligned} \tag{1}$$

where $\mathbf{\Pi}_j$ denotes the j^{th} column of $\mathbf{\Pi}$, $\text{Diag}(\mathbf{\Pi}_j)$ denotes a diagonal matrix with $\mathbf{\Pi}_j$ along the diagonal, $\alpha = d/(m\epsilon^2)$, $\alpha_j = d/(\langle \mathbf{1}, \mathbf{\Pi}_j \rangle \epsilon^2)$, $\gamma_j = \langle \mathbf{1}, \mathbf{\Pi}_j \rangle / m$, \mathcal{S} is the set of all matrices whose columns all have unit ℓ_2 norm³ and $\epsilon > 0$ is a prescribed precision error. Roughly speaking, $R(\mathbf{Z}_\theta)$, known as the *expansion* term, captures the dimension (or the volume) of the space spanned by \mathbf{Z}_θ while $R_c(\mathbf{Z}_\theta, \mathbf{\Pi})$, or the *compression* term, measures the sum of the dimensions/volumes of the data from each class. From an information-theoretic point of view, $R(\mathbf{Z}_\theta)$ estimates the coding rate, or the number of binary bits required to encode \mathbf{Z}_θ , through ϵ -ball packing [13]. The terms are called expansion and compression terms respectively, since by maximizing ΔR , the first coding rate term is maximized, which seeks to expand the overall volume of the embedded features, while the second coding rate term is minimized, which seeks to compress the volume of the embedded features from each class.

By assessing the MCR² objective (1), one can already observe a potential drawback of MCR² for optimization. In particular, note that each log det term requires $O(\min\{d^3, m^3\})$ operations to compute (and similarly to back-propagate through). While $d \ll D$ can often be made reasonably small for many high-dimensional datasets which have an underlying low-dimensional structure, $R_c(\mathbf{Z}_\theta, \mathbf{\Pi})$ in particular is still often expensive to compute because it requires k computations of log det. This severely limits the application of MCR² for datasets with even moderate numbers, say hundreds, of classes as the objective becomes computationally infeasible on common machines.

2.2. Variational Forms of Spectral Functions

To avoid this computational bottleneck, here we propose instead a formulation which takes advantage of variational

²Note that if the labels are known exactly then the entries of $\mathbf{\Pi}$ are binary $\{0, 1\}$ with each row of $\mathbf{\Pi}$ summing to one. Notice that our notation of $\mathbf{\Pi}$ is slightly different from that adopted in [21]. Our choice is more compact for optimization purposes.

³Note the constraint that \mathbf{Z}_θ has unit norm columns is often achieved by simply having the final operation of the network f_θ be a normalization.

forms of spectral functions of a matrix. Specifically, for a given positive semi-definite (PSD) matrix M and any scalar $c \geq 0$, note that

$$\log \det(\mathbf{I} + cM) = \sum_{i=1}^r \log(1 + c\sigma_i(M)), \quad (2)$$

where r denotes the rank of M and $\sigma_i(M)$ is the i^{th} singular value of M . Here note that $\log(1 + c\sigma)$ is a non-decreasing, concave function of σ , so we can exploit known variational forms of spectral functions [6, 16]. In particular note the following result in [16].

Theorem 2.1 (Adapted from [16]) *For any matrix X , let r denote the rank of X , let $\sigma_i(X)$ denote the i^{th} singular value of X , and define*

$$H(X) = \sum_{i=1}^r h(\sigma_i(X)).$$

for some function h . If h is a concave, non-decreasing function on $[0, \infty)$ with $h(0) = 0$, then the following holds

$$H(X) = \min_{U, V: UV^T = X} \sum_i h(\|U_i\|_2 \|V_i\|_2),$$

where (U_i, V_i) denotes the i^{th} columns of (U, V) . Note also that (U, V) can have an arbitrary number of columns ($\geq r$) provided $UV^T = X$.

3. Proposed Formulation

Having introduced the above background material, we now describe our proposed approach. In particular, note that Theorem 2.1 immediately gives the following result as a proposition.

Proposition 3.1 *Let M be any real positive semi-definite matrix and let $c \geq 0$ be any non-negative scalar. Then the following holds:*

$$-\log \det(\mathbf{I} + cM) = \max_{U: UV^T = M} -\sum_i \log(1 + c\|U_i\|_2^2). \quad (3)$$

Further, if $\bar{U}S\bar{U}^T = M$ is a SVD of M then $U^* = \bar{U}S^{1/2}$ is a solution to the above problem.

Proof. First recall the basic fact that for any function $\psi(x)$ one has $-\min_x \psi(x) = \max_x -\psi(x)$. Additionally, recall (2) and note that the function $h(x) = \log(1 + cx)$ satisfies the conditions required for h in Theorem 2.1. These facts give the following:

$$-\log \det(\mathbf{I} + cM) = \max_{U, V: UV^T = M} -\sum_i \log(1 + c\|U_i\|_2 \|V_i\|_2). \quad (4)$$

Further, note (4) implies that $-\log \det(\mathbf{I} + cM) \geq -\sum_i \log(1 + c\|U_i\|_2^2)$ for all U such that $UU^T = M$ since we have simply added the constraint $U = V$. The result is completed by noting that for the choice of $U = U^*$ the maximum can be attained since $\log(1 + c\|U_i^*\|_2^2) = \log(1 + c\sigma_i(M))$, $\forall i \in [r]$. ■

3.1. Variational Formulation of MCR²

Using Proposition 3.1, we develop our formulation by replacing the $-\log \det$ terms in R_c of (1) by the above variational form. In particular, note that for each class j we can eliminate the associated $-\log \det$ term in R_c of (1) by introducing an additional matrix $U^{(j)}$ subject to the constraint that $Z_\theta \text{Diag}(\Pi_j) Z_\theta^T = U^{(j)}(U^{(j)})^T$. Further, due to the fact that each row of Π sums to one, we also have $\sum_{j=1}^k Z_\theta \text{Diag}(\Pi_j) Z_\theta^T = Z_\theta Z_\theta^T$. As such, the variational form in Proposition 3.1 gives that the original MCR² objective in (1) is equivalent to the following constrained variational form:

$$\begin{aligned} \max_{\theta} \Delta R(Z_\theta) = & \max_{\theta, \{U^{(j)}\}_{j=1}^k} \frac{1}{2} \log \det \left(\mathbf{I} + \alpha \sum_{j=1}^k U^{(j)}(U^{(j)})^T \right) \\ & - \sum_{j=1}^k \frac{\gamma_j}{2} \sum_i \log \left(1 + \alpha_j \|U_i^{(j)}\|_2^2 \right) \\ \text{s.t. } \forall j, U^{(j)}(U^{(j)})^T = & Z_\theta \text{Diag}(\Pi_j) Z_\theta^T \text{ and } Z_\theta \in \mathcal{S}. \end{aligned} \quad (5)$$

From this form, we now reparameterize the $U^{(j)}$ matrices as $U^{(j)} = \Gamma \text{Diag}(A_j)^{1/2}$ where $\Gamma \in \mathbb{R}^{d \times q} \cap \mathcal{S}$ is a dictionary with unit norm columns, and $A_j \in \mathbb{R}_+^q$ is a (non-negative) encoding vector. Now let $A \in \mathbb{R}_+^{q \times k} = [A_1, \dots, A_k]$ be a matrix of the concatenated encoding vectors and note that we trivially have $\Gamma \text{Diag}(A_j) \Gamma^T = U^{(j)}(U^{(j)})^T$ and $\|U_i^{(j)}\|_2^2 = A_{i,j}$, which gives another equivalent formulation for the MCR² objective, provided the number of dictionary elements q is sufficiently large so that each optimal $U^{(j)}$ matrix in (5) can be encoded by Γ (i.e., each column of $U^{(j)}$ must be a column of Γ within a scaling factor):

$$\begin{aligned} \max_{\theta} \Delta R(Z_\theta) = & \max_{\theta, \Gamma \in \mathbb{R}^{d \times q} \cap \mathcal{S}, A \in \mathbb{R}_+^{q \times k}} \frac{1}{2} \log \det \left(\mathbf{I} + \alpha \sum_{j=1}^k \Gamma \text{Diag}(A_j) \Gamma^T \right) \\ & - \sum_{j=1}^k \frac{\gamma_j}{2} \sum_{l=1}^q \log(1 + \alpha_j A_{l,j}) \\ \text{s.t. } \forall j, \Gamma \text{Diag}(A_j) \Gamma^T = & Z_\theta \text{Diag}(\Pi_j) Z_\theta^T \text{ and } Z_\theta \in \mathcal{S}. \end{aligned} \quad (6)$$

Finally, we relax the strict equality constraints $\Gamma \text{Diag}(\mathbf{A}_j) \Gamma^\top = \mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$ with ℓ_2 penalties $\frac{1}{\gamma_j} \|\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top - \Gamma \text{Diag}(\mathbf{A}_j) \Gamma^\top\|_F^2$ to arrive at our final proposed formulation, which we call V-MCR²:

$$\max_{\theta, \Gamma \in \mathbb{R}^{d \times q} \cap \mathcal{S}, \mathbf{A} \in \mathbb{R}^{q \times k}} R^v(\Gamma, \mathbf{A}) - R_c^v(\mathbf{A}) - \frac{\mu}{2m} M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$$

where $R^v(\Gamma, \mathbf{A}) = \frac{1}{2} \log \det \left(\mathbf{I} + \alpha \sum_{j=1}^k \Gamma \text{Diag}(\mathbf{A}_j) \Gamma^\top \right)$,

$$R_c^v(\mathbf{A}) = \sum_{j=1}^k \frac{\gamma_j}{2} \sum_{l=1}^q \log(1 + \alpha_j \mathbf{A}_{l,j}),$$

$$M(\mathbf{Z}_\theta, \Gamma, \mathbf{A}) = \sum_{j=1}^k \frac{1}{\gamma_j} \left\| \mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top - \Gamma \text{Diag}(\mathbf{A}_j) \Gamma^\top \right\|_F^2, \quad (7)$$

such that $\mathbf{Z}_\theta \in \mathcal{S}$. Regularization parameter $\mu > 0$ weights how strictly the equality constraints should be approximated, and the $\frac{1}{\gamma_j}$ terms roughly ensure class balance (recall, $\gamma_j = \langle \mathbf{1}, \mathbf{\Pi}_j \rangle / m$). From this reformulation, we have significantly reduced the complexity of evaluating the objective function. The $\log \det$ terms in R_c which take $O(k \min\{d^3, m^3\})$ time to evaluate is now replaced by $O(qk)$ for the sum of the $\log(1 + \alpha_j \mathbf{A}_{l,j})$ terms along with the cost of computing the M term which scales as $O(kd^2)$.

3.2. Interpretation of the Variational Form

Besides the above computational advantages, we also discuss a few additional aspects of our formulation below.

Sparsifying dictionary learning interpretation. Notice that the above variational reformulation takes on a natural interpretation as learning a sparsifying dictionary: it essentially “parameterizes” the subspaces spanned by each class with a common shared dictionary Γ . Every class then selects a “sparse” number of eigenbases, with \mathbf{A}_j , from this dictionary and forms its estimate of the sample covariance within the subspace. Notice that the scalar log terms in $R_c^v(\mathbf{A})$ are precisely nonconvex sparsity promoting measures adopted in early studies of sparse representation [14, 15]. The sparsity in the (spectral) bases in terms of \mathbf{A}_j precisely corresponds to the subspace spanned by each class being low-dimensional or low-rank.

Penalty function method and other options. Notice that in our formulation (7), the equality constraint in (6) is enforced through a penalty function M . As the penalty weight μ increases to infinity and the dictionary Γ is sufficiently large⁴, the formulation becomes exactly equivalent to the original formulations (6) and (1). Of course, to deal with the equality constraint in (6) more precisely, one may also

⁴In the worst case the model become equivalent when Γ is large enough to contain a concatenation of the singular vectors of each class $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$.

consider adopting more advanced methods such as the augmented Lagrangian multiplier method to incorporate the equality constraint [7, 18], which we leave for future work. However, as we discuss next, by relaxing the strict equality constraint we also gain a potential advantage when the latent representation contains noise.

Low-rank LASSO interpretation. Notice that the sparse/low-rank promoting term $R_c^v(\mathbf{A})$ and the quadratic penalty term M together resemble the classic LASSO method for recovering a sparse solution from noisy measurements [20]. The only difference here is that we are seeking a sparse solution in the spectrum of a covariance matrix – hence seeking a low-rank solution for the covariance. So to some extent, one may consider the variational form as a “low-rank LASSO.” We have noticed a nice side benefit of this LASSO-type formulation: empirically it seems to lead to better solutions than solving the original MCR² objective (see experimental results in Section 5). Part of the reason is likely because the LASSO type relaxation introduced by the variational form finds a solution that is more stable to small noise in the data or deviation from an ideal low-dimensional linear subspace.

3.3. Optimization Strategy

3.3.1 Alternating Maximization

To optimize (7), we adopt an alternating maximization strategy [3] between the variation parameters (Γ, \mathbf{A}) and the network parameters (θ) . At each iteration, we first optimize Γ and \mathbf{A} by taking one step of a proximal gradient ascent update which consists of taking a gradient ascent step on the relevant part of variational loss $\Delta R^v - \frac{\mu}{2m} M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$, followed by normalizing the columns of Γ to have unit ℓ_2 norm and thresholding the negative entries of the updated \mathbf{A} matrix to 0 (i.e., applying the ReLU function). To ensure stability of the gradient-based method, we inversely scale the learning rate of Γ and \mathbf{A} by upper-bounds of the Lipschitz constants of the gradients, $\frac{1}{L_\Gamma}$ and $\frac{1}{L_{\mathbf{A}}}$, respectively. (See Appendix for our derivation for the bounds.) Next, the matrix approximation term $M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$ is recomputed using the updated Γ and \mathbf{A} , and we then update the network parameters, θ , by taking a gradient step on the relevant part of the variation loss, $\nabla_\theta M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$.

In addition, note that from the variational form in Proposition 3.1 we know that the optimal variational parameters (Γ, \mathbf{A}) should be closely related to the singular values and vectors of $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$ for each class j (the relationship becomes exact for large values of μ). We exploit this fact to initialize the variational parameters and to make periodic ‘approximately closed-form’ updates to the variational parameters. We call this procedure *latching*, which we describe in more detail in the next Section 3.3.2. We summa-

size our overall training process in Algorithm 1.⁵

Algorithm 1 Variational MCR² Training

```

1: Input: data  $\mathbf{X}$ , labels  $\mathbf{Y}$ , featurizer  $f_\theta(\cdot)$ ,
   latch-freq, step sizes  $(\nu_\theta, \nu_\Gamma, \nu_{\mathbf{A}})$ 
2: Initialize  $\mathbf{A}, \Gamma \leftarrow \text{latching}(\mathbf{X}, \mathbf{Y}, f_\theta)$ 
3: for iter = 0, 1, ...,  $n - 1$  do
4:   Get  $\mathbf{Z}_\theta = f_\theta(\mathbf{X})$  and membership matrices  $\mathbf{\Pi}$ 
5:   Get  $\ell_{\text{V-MCR}^2}(\mathbf{Z}_\theta, \Gamma, \mathbf{A}) = (7)$ 
6:   Compute  $L_{\mathbf{A}}, L_\Gamma$  (see Appendix)
7:    $\Gamma \leftarrow \Gamma + \frac{\nu_\Gamma}{L_\Gamma} \nabla_\Gamma \ell_{\text{V-MCR}^2}(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$ 
8:    $\mathbf{A} \leftarrow \mathbf{A} + \frac{\nu_{\mathbf{A}}}{L_{\mathbf{A}}} \nabla_{\mathbf{A}} \ell_{\text{V-MCR}^2}(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$ 
9:   Project  $\mathbf{A} \leftarrow \text{ReLU}(\mathbf{A})$ 
10:  Project  $\Gamma_l \leftarrow \frac{1}{\|\Gamma_l\|_2} \Gamma_l \quad \forall l \in [q]$ 
11:  Recompute  $M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$ 
12:   $\theta \leftarrow \theta - \nu_\theta \nabla_\theta (M(\mathbf{Z}_\theta, \Gamma, \mathbf{A}))$ 
13:  if iter mod latch-freq = 0 then
14:     $\mathbf{A}, \Gamma \leftarrow \text{latching}(\mathbf{X}, \mathbf{Y}, f_\theta)$ 
15:  end if
16: end for
17: return  $f_\theta$ 

```

Algorithm 2 Latching

```

Input: data  $\mathbf{X}$ , labels  $\mathbf{Y}$ , featurizer  $f_\theta(\cdot)$ 
Get  $\mathbf{Z}_\theta = f_\theta(\mathbf{X}) \in \mathbb{R}^{d \times m}$  and membership  $\mathbf{\Pi} \in \mathbb{R}^{m \times k}$ 
 $\mathbf{A} \leftarrow \mathbf{0} \in \mathbb{R}^{q \times k}$  (assume  $q$  is divisible by  $k$ )
 $\Gamma \leftarrow \mathbf{0} \in \mathbb{R}^{d \times q}$ 
for  $j = 1, \dots, k$  do
  Get  $\mathbf{U} \text{Diag}(\boldsymbol{\sigma}) \mathbf{V}^\top = \text{SVD}(\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top)$ 
   $s \leftarrow q/k$ 
   $\Gamma[:, (j-1)*s : j*s] = \mathbf{U}[:, 0 : s]$  % python indexing
   $\mathbf{A}[(j-1)*s : j*s, j] = \boldsymbol{\sigma}[0 : s]$  % python indexing
end for
return  $\mathbf{A}, \Gamma$ 

```

3.3.2 Latching

In order to optimize the variational MCR² objective, the dictionary Γ and \mathbf{A} must maximize $\Delta R(\mathbf{Z}_\theta) = R^v(\Gamma, \mathbf{A}) - R_c^v(\mathbf{A})$ whilst minimizing the ℓ_2 regularization term $M(\mathbf{Z}_\theta, \Gamma, \mathbf{A})$. This trade-off is controlled by the regularization constant μ . In practice, when μ is too large, each gradient step does not allow for $\Gamma \text{Diag}(\mathbf{A}_j) \Gamma^\top$ to stray too far away from $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$, which can result in slow convergence. We observe that the following procedure improves convergence in practice. Note that the varia-

⁵Note that for clarity we describe the full procedure for gradient ascent, but in our experiments stochastic gradient ascent is implemented.

tional form is maximized⁶ when Γ and \mathbf{A} are derived from the SVDs of $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$, as given in Proposition 3.1. This gives a means to periodically reinitialize the variational parameters (Γ, \mathbf{A}) based on the SVDs of $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$, which we refer to as *latching* as described in detail in Algorithm 2. This latching step can be viewed as taking an (approximate) full-maximization step w.r.t. the variational parameters (as opposed to a proximal gradient descent step) based on the closed-form solution provided in Proposition 3.1. This will be an exact maximization step as μ becomes large. In short, given a dictionary with q columns, we initialize the dictionary as the concatenation of the top q/k singular vectors of $\mathbf{Z}_\theta \text{Diag}(\mathbf{\Pi}_j) \mathbf{Z}_\theta^\top$ for each class $j \in [k]$. Similarly, the columns of \mathbf{A} are initialized as the corresponding singular values. Though latching is in itself an expensive procedure, requiring one to compute the SVD of a matrix k times, it is optional (though we notice a benefit in practice) and can be done relatively infrequently throughout training with a proper choice of the hyper-parameter `latch-freq`. As a result, the amortized cost of latching becomes insignificant.

4. Experimental Setup

We compare ΔR , wall-clock time, and accuracy of models trained with the original MCR² objective and the variational MCR² objective on MNIST [12], CIFAR-10 [10], CIFAR-100 [10], and Tiny ImageNet [11] (with 200 classes) datasets. We also compare the performance to cross-entropy (CE) training as a benchmark of the correctness of the learned representations (for classification). The high-level goal of these experiments is to show that 1) the variational MCR² objective is feasible for datasets where the original MCR² objective is computationally expensive (or impossible) to train (such as CIFAR-100 and Tiny ImageNet), 2) show that training on the variational MCR² maximizes the true ΔR objective and obtains the desired subspace-like representations.

4.1. Hyperparameters

For fair comparison across training objectives (original MCR², variational MCR²), we use a learning rate of 10^{-3} for the network optimizer and the same batch size. For CE, we use the same batch size, but a larger learning rate of 10^{-2} . The batch size is 1000 for MNIST and CIFAR-10, and 2000 for CIFAR-100 and Tiny ImageNet. The network is optimized using stochastic gradient descent for all objectives. For the precision error ϵ of the MCR² objectives, we use $\epsilon^2 = 0.5$ for all datasets. ϵ and batch size are consistent with the experimental settings in the original MCR² work [21] for MNIST and CIFAR-10. For Variational MCR², the regularization constant $\mu = 1$ and initial

⁶The maximization is exact as μ becomes large, but a good approximation otherwise.

learning rates $\nu_\Gamma = 5$, $\nu_A = 5$ across all experiments. We perform latching every 50 epochs for all experiments. See Appendix for precise details. The dictionary size q and feature dimension d varies across the datasets. For MNIST and CIFAR-10 we use $d = 128$ and $q = 20 \cdot k$, and for CIFAR-100 and Tiny ImageNet we use $d = 500$ and $q = 10 \cdot k$.

4.2. Nearest Subspace Classifier

MCR² is a loss over the featurizer f_θ . To classify the test data, we use the nearest-subspace classifier similar to the original MCR² work [21]. As shown by [21], at the global optima of MCR², the representations of each class lie on low-dimensional subspaces that are orthogonal to each other. Yu, et al. [21] also empirically observe this property for networks trained by SGD. Assuming that the learned representations satisfy this property, given a test datapoint, we can simply identify the closest subspace for the final classification. Formally, given a test sample $\mathbf{z}_{test} = f_\theta(\mathbf{x}_{test})$, the predicted label is given as

$$y = \arg \min_{j \in \{1, \dots, k\}} \left\| (\mathbf{I} - \mathbf{V}^{(j)}(\mathbf{V}^{(j)})^\top) \mathbf{z}_{test} \right\|_2^2, \quad (8)$$

where $\mathbf{V}^{(j)}$ is a matrix of the top $\lfloor \frac{d}{k} \rfloor$ principal components of $\mathbf{Z}_\theta \text{Diag}(\boldsymbol{\Pi}_j) \mathbf{Z}_\theta^\top$ with $\mathbf{Z}_\theta = f_\theta(\mathbf{X})$ being the embedding of the training data \mathbf{X} .

5. Experimental Results

We discuss the performance of variational MCR² below. Performance is measured by 1) training speed, 2) the true ΔR value the model reaches over the training data, and 3) the classification accuracy over the test set. Note that we are more interested in comparing the accuracy and training efficiency between MCR² and V-MCR², over trying to reach state of the art results in this paper.

5.1. Computational Efficiency

Across datasets, we compare the wall-clock time to train one epoch using the MCR² (1) and V-MCR² (7) formulations. The batch size is set to be the same for both models, and all our experiments are performed using PyTorch 1.9.0 and Python 3.8.11 on Nvidia A100-SXM4 GPUs with 40GB of CUDA memory for fair comparison. As shown in Table 1, V-MCR² training completes approximately 5× faster on CIFAR-100 and 12× faster on Tiny ImageNet. Even for datasets containing a small number of classes, i.e., MNIST and CIFAR-10, we observe a 1.5 – 2× speedup. Note that the overhead for the original MCR² model escalates significantly as the number of classes increases, so we expect even greater improvements in training efficiency with datasets with more classes. Even for Tiny ImageNet, training until convergence using MCR² becomes nearly impractical, while it is easily handled by V-MCR².

Dataset	MCR ²	V-MCR ²
MNIST	11.56	6.29
CIFAR-10	33.06	20.71
CIFAR-100	157.45	31.14
Tiny ImageNet	527.85	44.23

Table 1. **Wall-Clock Time per Epoch (in seconds).** We compare the wall-clock time to complete one epoch of training. MNIST and CIFAR-10 use a batch size of 1000. For CIFAR-100 and Tiny ImageNet, we use a batch size of 2000.

Additionally, we compare the true $\Delta R(\mathbf{Z}_\theta)$ (i.e., computing the original MCR² objective with the iterates of the V-MCR² model) over training epochs and observe that both models 1) take approximately the same number of training epochs to converge and 2) reach approximately the same final $\Delta R(\mathbf{Z}_\theta)$ objective value at convergence. Thus, V-MCR² does not require additional epochs to obtain a good solution which might offset the increased efficiency per epoch. As shown in Figure 1, on MNIST and CIFAR-10, V-MCR² and MCR² follow a similar training loss trajectory across epochs. For CIFAR-100, we observe that the convergence rate depends on the dimension of the features/representations, d . For a feature dimension of $d = 100$ for CIFAR-100, we observed similar number of epochs to convergence for MCR² and V-MCR². With a feature dimension of $d = 500$, we observe across 5 seeds that training with the original MCR² leads to a rapid convergence to a poor local optima where the expansion term $R(\mathbf{Z})$ increases rapidly but the compression term $R_c(\mathbf{Z})$ remains the same. On the other hand, although V-MCR² requires more epochs to converge, we observe the standard expected behavior where $R(\mathbf{Z})$ increases and $R_c(\mathbf{Z})$ decreases and the final solution is of much higher quality (see 5.2). Experiments on Tiny ImageNet show similar behaviors.

5.2. V-MCR² Obtains Better Representations

There are two properties of the representation that we aim to attain by optimizing the MCR² objective. We say that a representation of the training data is of ‘high quality’ if points from different classes lie on separate, orthogonal subspaces, and the union of these subspaces span as many dimensions as possible. In particular, the orthogonal property is important in order to classify the points using the nearest subspace algorithm in Section 4.2. To check the orthogonality of subspaces learned by f_θ , we report the inner product between every pair of training points as a heatmap in Figure 2. Namely, we sort the columns of \mathbf{Z}_θ by class and compute $|\mathbf{Z}_\theta^\top \mathbf{Z}_\theta|$. Ideally, we want $|\mathbf{Z}_\theta^\top \mathbf{Z}_\theta|$ to have a block diagonal structure, with $(\mathbf{Z}_\theta)_i^\top (\mathbf{Z}_\theta)_j \approx 0$ for i, j notating points from different classes.

Figure 2 shows heatmaps for MNIST, CIFAR-10, and CIFAR-100 after 2000 training epochs, when ΔR has con-

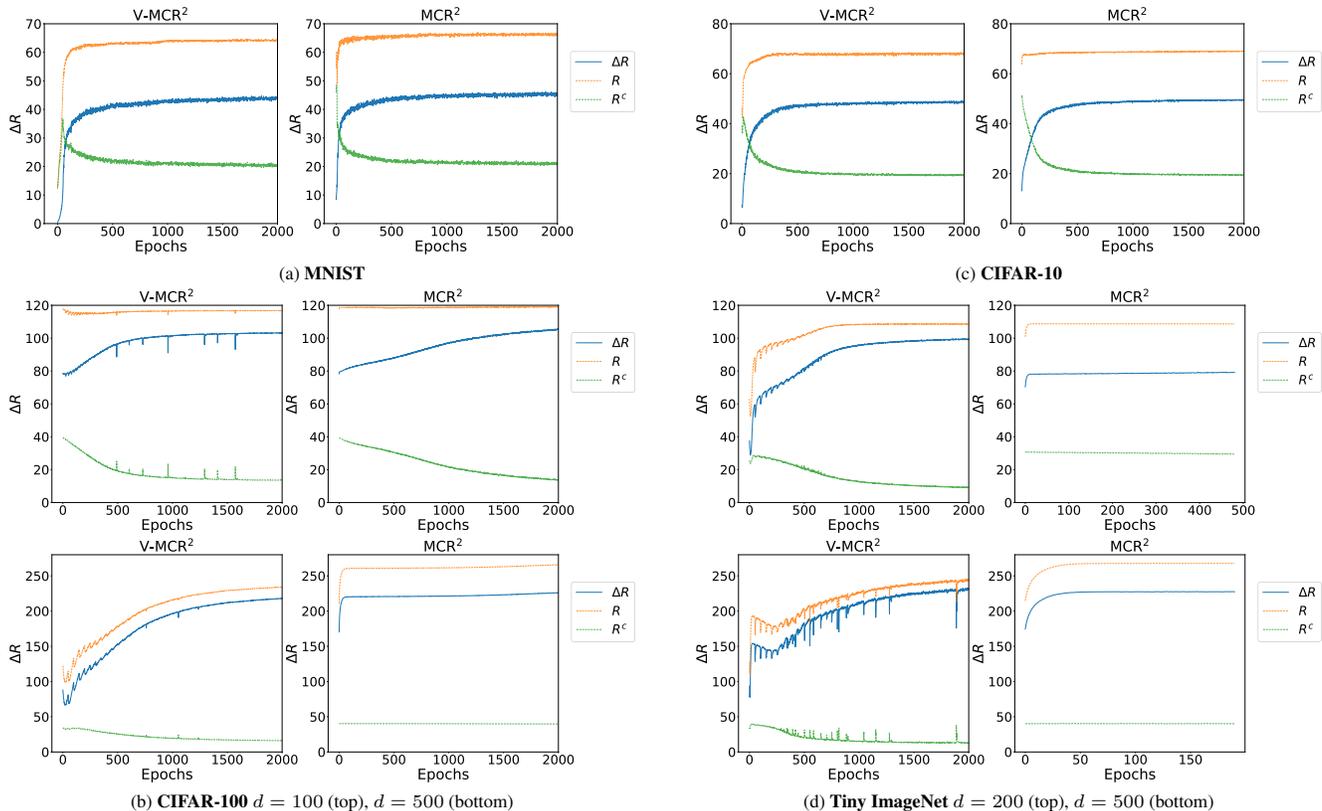


Figure 1. **Convergence of training ΔR .** We compare the $\Delta R(Z_\theta)$ of the training data over epochs for V-MCR² and MCR². For both V-MCR² and MCR², the network is optimized by stochastic gradient descent with a learning rate of 10^{-3} . All training runs are 2000 epochs, excluding the MCR² runs for Tiny ImageNet, which we stop at 500 and 200 for $d = 200$ and $d = 500$ respectively, due to significant computational cost from the high number of classes. Also, note that for V-MCR², we often observe small undulations in the training loss due to the regular reinitialization by latching.

verged for both MCR² and V-MCR² (Figure 1). For MNIST and CIFAR-10, note that both MCR² and V-MCR² obtain a block diagonal structure. However, on CIFAR-100 and Tiny Imagenet, we observe no block diagonal structure after MCR² training, whereas we see a clear block diagonal structure for V-MCR². These findings suggest that V-MCR² training is more robust to avoiding poor local minima than training on the original MCR² model, particularly as the number of classes increase. We leave a rigorous study of these phenomena for a future work.

5.3. Performance on Classification Tasks

In Table 2, we present test accuracies on the four aforementioned datasets when trained under original MCR² and V-MCR² objectives. In addition, we also train a separate model for each dataset by using CE loss as a reference classifier and report its test accuracy. Notice that the goal of our study here is *not* about achieving the best possible classification accuracy on these datasets – the training procedure and architectures used here are not optimal for that pur-

pose⁷. Instead, we make fair comparison of all methods on the same networks and datasets to justify the computational efficiency and effectiveness of the proposed method. To ensure fairness, we initialize this reference model with the same architecture along with other hyperparameters⁸ as in V-MCR² experiments and attach a final linear classifier with output dimension corresponding to the number of classes.

As shown in Table 2, when trained on datasets with a small number of classes, all three training objectives can reach competitive classification performance. We can observe that both MCR² and V-MCR² objectives are comparable to CE in these small-scale datasets. Again on CIFAR-100 we observe that the poor local minima obtained by training the original MCR² objective results in a poor test accuracy, while the V-MCR² model achieves comparable performance to training the same network with CE. Due to

⁷We use simple training practices (*e.g.* input data downsampling, minimal data augmentation and training from scratch), leading to about 15% and 30% difference in performance on CIFAR-100 and Tiny ImageNet as reported in [22] from much more carefully engineered training recipes.

⁸For experiments using CE, we use a higher learning rate of 10^{-2} to improve convergence for CE training.

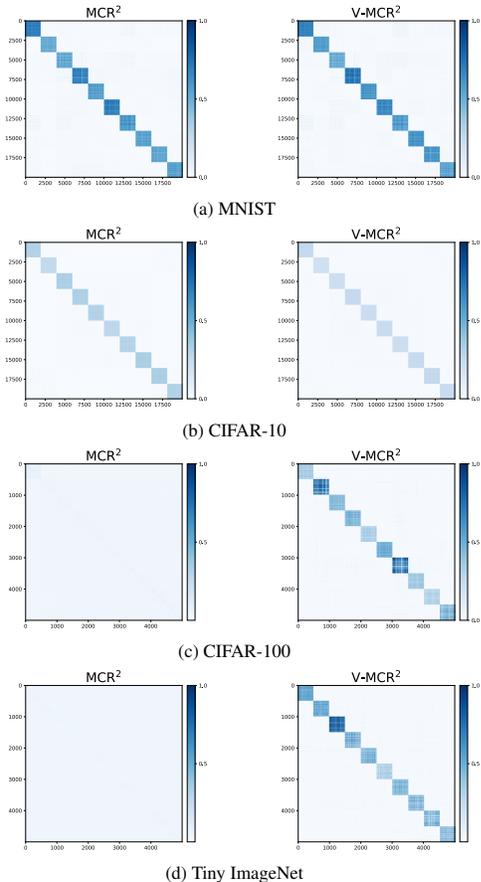


Figure 2. **Inner product of representations.** We plot the heatmap of $|Z_\theta^\top Z_\theta|$ where Z_θ are the representations of the training data, ordered by the class they belong to. For CIFAR-100 and Tiny ImageNet, 10 classes are randomly chosen. If the classes lie on low-dimensional orthogonal subspaces, we expect to see a block diagonal structure.

limited resources and heavy computational requirements of MCR^2 training with a large number of classes, we can only report the result on Tiny ImageNet by optimizing MCR^2 after 200 epochs while the other two objectives complete the full training session of 2000 epochs.

6. Conclusion

Building on Yu et al. [21], we propose an alternative cost-efficient formulation of the MCR^2 objective that is scalable to datasets with a large number of classes. Namely, for CIFAR-100 and Tiny ImageNet, we observed a $5\times$ and $12\times$ speedup per training epoch, respectively. The gain would be even more significant as the number of classes increases. Additionally, we show that we do not make any compromises when it comes to performance by using this approximate, variational formulation. In fact, in all datasets we tested, not only does $V-MCR^2$ reach similar ΔR values as MCR^2 , the learned representations of $V-MCR^2$ were just

Dataset	Objective	Training ΔR	Test Accuracy
MNIST	MCR^2	44.6429	0.9785
	$V-MCR^2$	44.2117	0.9788
	CE	-	0.9738
CIFAR-10	MCR^2	49.40	0.8956
	$V-MCR^2$	48.43	0.8997
	CE	-	0.8665
CIFAR-100	MCR^2	226.0519	0.2421
	$V-MCR^2$	218.0185	0.5872
	CE	-	0.5840
Tiny ImageNet 200	MCR^2	227.6468	0.1319
	$V-MCR^2$	231.1538	0.2665
	CE	-	0.1907

Table 2. **Comparison of classification performance.** We evaluate the training ΔR and test accuracy of CE, MCR^2 , and $V-MCR^2$ after 2000 training epochs for MNIST, CIFAR-10, and CIFAR-100. For Tiny ImageNet, we report the results after 200 epochs for MCR^2 due to very slow training, though note that ΔR had already converged (Fig. 1d). For $V-MCR^2$ and CE, results are after 2000 epochs. For CIFAR-100 and Tiny ImageNet, we report results for $d = 500$. See Appendix for results for other choices of d .

as good, oftentimes better, than those learned by MCR^2 for the same number of training epochs. On a related note, we observe an interesting phenomenon that directly optimizing the MCR^2 objective is not only slow, but often completely fails to learn the desired orthogonal subspace structures when the number of classes increases even after ΔR has successfully converged. In our experiments, we found that $V-MCR^2$ is surprisingly not prone to this problem, suggesting that $V-MCR^2$ allows for better control over learning the desired representations. By these means, we find that $V-MCR^2$ is a very promising adaptation of MCR^2 and we leave a rigorous comparison of the representations learned by MCR^2 and $V-MCR^2$ for future work.

Acknowledgements This work was partially supported by NSF Fellowship DGE2139757, NSF grants 1704458 and 2031985, and the Northrop Grumman Mission Systems Research in Applications for Learning Machines (REALM) initiative. Yi Ma was funded by ONR grants N00014-20-1-2002 and N00014-22-1-2102, the joint Simons Foundation-NSF DMS grant 2031899.

References

- [1] Which shortcut cues will dnns choose? a study from the parameter-space perspective. *arXiv preprint arXiv:2110.03095*, 2021. 1
- [2] Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation, and self-distillation in deep learning. <https://arxiv.org/abs/2012.09816>, 2021. 1
- [3] J. Eckstein. Augmented Lagrangian and alternating direction methods for convex optimization: A tutorial and some illustrative computational results. *RUTCOR Technical Report*, 2012. 4
- [4] M. Fazel, H. Hindi, and S.P. Boyd. Log-det heuristic for matrix rank minimization with applications to hankel and euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003.*, volume 3, pages 2156–2162 vol.3, 2003. 2
- [5] R Geirhos, J.-H. Jacobsen, C Michaelis, R Zemel, W Brendel, Bethge M, and F. A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2020. 1
- [6] Paris Giampouras, René Vidal, Athanasios Rontogiannis, and Benjamin Haeffele. A novel variational form of the Schatten- p quasi-norm. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [7] Magnus R Hestenes. Multiplier and gradient methods. *Journal of optimization theory and applications*, 4(5):303–320, 1969. 4
- [8] Like Hui, Mikhail Belkin, and Preetum Nakkiran. Limitations of neural collapse for understanding generalization in deep learning. *arXiv preprint arXiv:2202.08384*, 2022. 1
- [9] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 2
- [10] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [11] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *Stanford cs231*, 2015. 5
- [12] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5
- [13] Yi Ma, Harm Derksen, Wei Hong, and John Wright. Segmentation of multivariate mixed data via lossy data coding and compression. *IEEE transactions on pattern analysis and machine intelligence*, 29(9):1546–1562, 2007. 2
- [14] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996. 4
- [15] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997. 4
- [16] Marcus Valtonen Ornhag, Carl Olsson, and Anders Heyden. Bilinear parameterization for differentiable rank-regularization. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020. 3
- [17] Vardan Pappayan, X.Y. Han, and David Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences of the United States of America*, 2020. 1
- [18] R Tyrell Rockafellar. The multiplier method of Hestenes and Powell applied to convex programming. *Journal of Optimization Theory and Applications*, 12(6):555–562, 1973. 4
- [19] Harshay Shah, Kaustav Tamuly, Aditi Raghunathan, Prateek Jain, and Praneeth Netrapalli. The pitfalls of simplicity bias in neural networks. *Advances in Neural Information Processing Systems*, 2020. 1
- [20] R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society B*, 58(1):267–288, 1996. 4
- [21] Yaodong Yu, Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. *NeurIPS*, 2020. 1, 2, 5, 6, 8, 10
- [22] Sukmin Yun, Jongjin Park, Kimin Lee, and Jinwoo Shin. Regularizing class-wise predictions via self-knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13876–13885, 2020. 7, 12