α-ReQ : Assessing representation quality by measuring eigenspectrum decay

Kumar Krishna Agrawal[†]
UC Berkeley
CA, USA

Arnab Kumar Mondal[†]
Mila & McGill University
Montréal, QC, Canada

Arna Ghosh[†]
Mila & McGill University
Montréal, QC, Canada

Blake A. Richards
Mila, Montreal Neurological Institute & McGill University
Montréal, QC, Canada
Learning in Machines and Brains Program, CIFAR
Toronto, ON, Canada

Abstract

Self-Supervised Learning (SSL) with large-scale unlabelled datasets enables learning useful representations for multiple downstream tasks. However, assessing the quality of such representations efficiently poses nontrivial challenges. Existing approaches train linear probes (with frozen features) to evaluate performance on a given task. This is expensive both *computationally*, since it requires retraining a new prediction head for each downstream task, and *statistically*, requires task-specific labels for multiple tasks. This poses a natural question, *how do we efficiently determine the "goodness" of representations learned with SSL across a wide range of potential downstream tasks?* In particular, a task-agnostic statistical measure of representation quality, that predicts generalization without explicit downstream task evaluation, would be highly desirable.

In this work, we analyze characteristics of learned representations $\mathbf{f}_\theta$, in well-trained neural networks with canonical architectures & across SSL objectives. We observe that the eigenspectrum of the empirical feature covariance $\mathrm{Cov}(\mathbf{f}_\theta)$ can be well approximated with the family of power-law distribution. We analytically and empirically (using multiple datasets, e.g. CIFAR, STL10, MIT67, ImageNet) demonstrate that the decay coefficient α serves as a measure of representation quality for tasks that are solvable with a linear readout, i.e. there exist well-defined intervals for α where models exhibit excellent downstream generalization. Furthermore, our experiments suggest that key design parameters in SSL algorithms, such as BarlowTwins [1], implicitly modulate the decay coefficient of the eigenspectrum (α). Notably, α depends only on the features themselves, and as such, using this measure for model selection in hyperparameter tuning for BarlowTwins allows searching for optimal hyperparameters with upto $\sim 15\times$ less compute.

1 Introduction

The recent success of self-supervised learning (SSL) has changed the landscape of deep learning significantly. With well-engineered architectures and training objectives, SSL models learn useful representations from large datasets without relying on any labels [1, 2, 3]. Despite this progress, quantifying the *representation quality* for models trained with SSL is still an open problem. The

[†] These authors contributed equally to this work

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

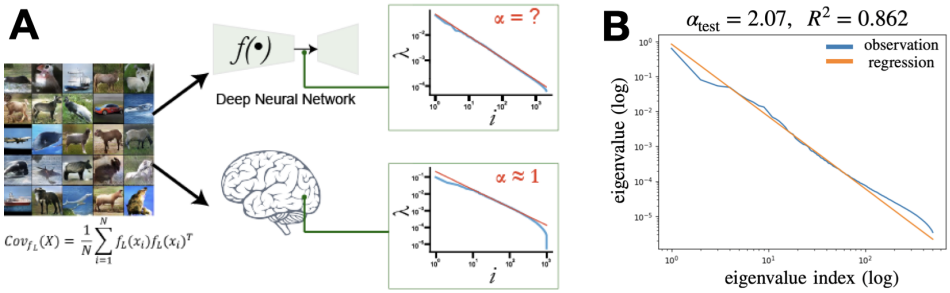

Figure 1: **(A)** One approach to evaluate representation quality is tracking the eigenspace of feature covariance matrix $\Sigma_n(\mathbf{f}) = 1/n \sum_{i=1}^{n} \mathbf{f}(x_i)\mathbf{f}(x_i)^\top$. Analysing populations of neural activation [6, 8], suggests that the eigenspectrum for $\Sigma_n(\mathbf{f})$ can be approximated by power-law, where $\lambda_i \propto i^{-\alpha}$. **(B)** For a ResNet-50 model pretrained on ImageNet, we extract activations from an intermediate layer. We plot the eigenvalue spectrum $\{\lambda_1, ..., \lambda_d\}$ for covariance matrix Σ_n (`resnet50(feats='block4')`) in `log-log` scale. Further, fitting a linear regressor we estimate the decay coefficient $\hat{\alpha}$.

most obvious (and common) solution is to assess performance of models using these representations on downstream tasks. However, if the goal is general representations that can be used across many domains, then this either requires a large investment of time and energy to be done well (in order to assess performance on many tasks and datasets). Alternatively, if models are assessed on only a small number of datasets and tasks, then it is hard to be confident in the assessment. Thus, we are left with a question: Can we assess the quality of learned representations without explicitly evaluating the performance on downstream tasks? To answer this question, we must formally define the "quality" of representations and subsequently examine statistical estimators that measure this property without needing downstream evaluation. Beyond theoretical interest, such a metric would be highly desirable for model selection and also useful for designing new SSL algorithms.

In search of such a metric, we turn our attention to one of the more efficient learning machines in existence – the mammalian brain. The hierarchical and distributed organization of neural circuits, especially in the cortex, provides neural representations that support a wide array of behaviours across many domains. For example, representations in primary visual cortex (V1) of mammalian brains are used by animals to support downstream behaviours ranging from object categorization to movement detection and motor control [4, 5]. This berth of downstream uses of V1 representations is desirable for artificial vision systems trained with SSL. Thus, understanding the properties of representations in V1 is a reasonable starting point for seeking a general metric of representation quality.

Recent breakthroughs in systems neuroscience enable large-scale recordings of neural activity. By recording and analyzing the response to visual stimuli, [6, 7] find that activations in the mouse and macaque monkey V1 exhibit a characteristic information geometric structure. In particular, these representations are high-dimensional, yet the amount of information encoded along the different principal directions varies significantly. Notably, this variance (computed by measuring the eigenspectrum of the empirical covariance matrix) is well-approximated by a power-law distribution with decay coefficient ≈ 1, i.e., the nth eigenvalue of the covariance matrix scales as $1/n$.

Motivated by these results, we explore the use of the decay rate of the empirical eigenspectrum to characterize representation *quality* in neural networks trained with SSL (Fig. 1a). Across diverse model architectures, pretraining objectives, and downstream classification tasks, we empirically observe an eigenspectrum decay in the representation covariance matrix that roughly follows a power-law distribution (see e.g. Fig. 1b). We also find that the coefficient of this decay, denoted by α, is informative of downstream generalization performance. Importantly, α can efficiently be calculated without any labels, and thereby, could be incorporated into existing SSL pipelines for efficient model selection on fixed compute budget. Our core contributions in this paper are:

1. α **as a potential metric** We observe that canonical pretrained architectures have representations which loosely conform to a power-law distribution in their covariance eigenspectrum. Under an assumption of a power-law distribution, we prove that the convergence rate and upper bound on generalization error are related to decay-coefficient (α) of the corresponding power-law distribution, with best values of α being neither too large nor too small.

2. α **as a label-free measure for representation quality**: We empirically validate our theoretical results by demonstrating a relationship between α and downstream generalization. In particular, we find that either too high or too low an α value implies poor generalization, both in-distribution and out-of-distribution. Generally, the best representations are those where α is in a range that is close to 1, as observed in V1 of the real brain. Furthermore, these results hold irrespective of the choice of network architecture or pretraining objective.
3. α **for model selection in SSL**: We establish α as a reliable metric for model selection in a specific case of SSL, Barlow Twins [1]. Notably, we show that α allows us to identify model hyper-parameters that lead to representations that generalize well without any labels, more so than the actual loss function used to train the network.

Altogether, our results show that α is a promising task/architecture/data agnostic metric for assessing representation quality in SSL. We publicly release our results and code [†].

2 Theoretical Framework to Assess Representations in SSL

We are interested in evaluating the quality of high dimensional representations learned by neural networks. Formally, we consider the overparameterized setting, with datasets $\mathcal{X} \subset \mathbb{R}^d$, and learned mappings $\mathbf{f} : \mathbb{R}^d \to \mathbb{R}^D$ such that $\mathbf{f}(\mathbf{x})$ is a vector of D-dimensional features.

We consider DNNs as our function approximators, where each architecture implicitly defines a function class $\mathcal{F} = \{\mathbf{f}_\theta : \theta \in \Theta\}$ where $\Theta \subset \mathbb{R}^p$ is the feasible set of model parameters (e.g bounded $\Theta, [-B, B]^p$ for some $B \in \mathbb{R}$). The search for good representations poses the following optimization problem: with a dataset $\mathcal{D}_{\text{pretrain}}$ from some data distribution $\mathbb{P}_{\text{pretrain}}$ (potentially with labels), search for *optimal* parameters θ^* such that with pretrain objective $\mathcal{L}_{\text{pretrain}}(\mathbf{f}, \mathcal{D})$

$$\theta^* = \arg\min_{\mathbf{f}_\theta \in \mathcal{F}} \mathcal{L}_{pretrain}(\mathbf{f}_\theta, D_{\text{pretrain}}) \tag{1}$$

The above optimization problem is usually non-convex, and often uses gradient based optimizer to find an approximate solution $\hat{\theta}$. Typically, to measure the quality of $f_{\hat{\theta}}$, researchers evaluate the quality of representations on a downstream task $\mathcal{D}_{\text{downstream}}$ with a linear readout using some metric $\mathcal{R}(\texttt{linear}(\mathbf{f}_{\hat{\theta}}), D_{\text{downstream}})$. A concrete example is pretraining a VGG-16 model ($\mathcal{F}$) on ImageNet dataset (D_{pretrain}) and evaluating the learned representations using classification accuracy ($\mathcal{R}$) by learning a linear classifier on the MIT67 dataset ($D_{\text{downstream}}$). For the rest of the paper, we denote feature maps as $\mathbf{f}_\theta(x) \in \mathbb{R}^D$ where $\mathbf{f}_\theta : \mathcal{X} \to \mathbb{R}^D$, and the readout network as $\mathbf{g}_\phi : \mathbb{R}^D \to \mathbb{R}^k$, where k is the target dimensionality. For simplicity of analysis, we consider linear readouts unless explicitly mentioned, i.e. $\mathbf{g}_\phi(x) = x^T\phi$.

2.1 Covariance estimation and eigenspectrum

For a parameterized function $\mathbf{f}_\theta : \mathcal{X} \to \mathbb{R}^D$ (assume centered), the (n-sample) empirical covariance matrix is defined as:

$$\Sigma_n(\mathbf{f}_\theta) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{f}_\theta(x_i)\mathbf{f}_\theta(x_i)^T$$

The eigenspectrum of $\Sigma_n(\mathbf{f}_\theta)$ informs us about the variance explained by each principal component of the space spanned by representations $\mathbf{f}_\theta(\mathbf{x})$. Using the spectral decomposition theorem on symmetric matrices, $\Sigma = U\Lambda U^T$, where Λ is a diagonal matrix with nonnegative entries, and U is a matrix whose columns are the eigenvectors of Σ. Without loss of generality we assume that $\lambda_1 \geq \lambda_2 ... \geq \lambda_m$, where $m = \min(n, D)$ is the rank of $\Sigma_n(\mathbf{f}_\theta)$.

2.2 Eigenspectrum Decay in Deep Representation Learning

Recent work in characterizing representation structure in canonical DNNs has demonstrated that the covariance eigenspectrum roughly follows a power-law [9, 10, 11]. Specifically, the eigenspectrum

[†] https://github.com/kumarkrishna/fastssl

of a covariance matrix follows a power-law distribution $PL(\alpha)$, or *zeta distribution* if for $\lambda_j \in [\lambda_{min}, \lambda_{max}]$, the eigenvalues λ_j are all nonnegative, and

$$\lambda_j \propto j^{-\alpha}$$

for some $\alpha > 0$. Here, α is the *slope* of the power law, and is referred to as the *coefficient of decay* of the eigenspectrum. Intuitively, small α (typically $\alpha \leq 1$) suggests a dense encoding, while a high α (rapid decay) corresponds to a sparse encoding.

Insights from High-Dimensional Linear Regression : Recent work in theoretical machine learning has connected bounds on generalization error for a linear regression problem to the eigenspectrum of the feature covariance matrix. Specifically, in the infinite-dimensional setting with $D \to \infty$, [12] studied the linear regression setting with Gaussian features, and proved that if the eigenspectrum follows a power-law distribution (up to polylogarithm factors), the min-norm solution provides good generalization performance iff $\alpha = 1$. The asymptotic regime of infinite width is an excellent framework to study theoretical properties of DNN representations. However, despite having a large number of parameters, practical DNNs always possess finite dimensional representations, making it important to investigate the implications of such results in finite width models. In particular, for the finite dimensional setting we try to answer: *Does α sufficiently larger or smaller than 1, still allow efficient learning and strong generalizability?*

To answer this question, we narrow our focus to gradient-based optimization techniques usually used to train DNNs. In particular, from the optimization perspective, Advani *et al.* showed that for deep linear regression in high dimensions, the time required for training and the steady-state generalization error are both $\mathcal{O}(\frac{1}{\lambda_{min}})$ [13]. A key difference from our work is that they assume the inputs are drawn from an isotropic Gaussian distribution. Instead, we investigate the generalization of linear regression on $\mathbf{f}_\theta(x)$, which has a power law structure in its covariates. We show that the training convergence time grows exponentially with α using the following theorem:

Theorem 2.1. *Let $\hat{y} = \mathbf{f}_\theta(x)^T \psi$ be an overparameterized linear regression problem where ψ is learned using gradient descent in order to optimize the training error, $\mathbb{E}_{x,y}[(y - \mathbf{f}_\theta(x)^T\psi)^2]$, where $(x, y) \sim \mathcal{D}_{train}$. If we assume a power-law distribution in the eigenspectrum of representations at $\mathbf{f}_\theta$, i.e. $\lambda_n = \frac{c}{n^\alpha} \quad \forall n \geq n^*$, where $n^* \in \{1, 2...N\}$, then the time required by gradient descent to minimize the training error, $T_{convergence} = \mathcal{O}(N^\alpha)$ where N is number of training samples.*

The outline of the proof builds on key results relating to gradient descent dynamics from [14]. We show that gradient descent updates, when ψ is initialized to 0, yield a recursive relation for $\psi(k)$, i.e. ψ after k update steps. Plugging this relation in the gradient formulation, we show that the update step length along the nth principal direction of $\mathbf{f}_\theta(x)$ shrinks exponentially with a decay rate proportional to λ_n. Therefore, the time to convergence in training is controlled by the smallest eigenvalue which, by design, follows the power law. In sum, Theorem 2.1 provides an explanation against arbitrarily large values of α.

Next we show that α can't be arbitrarily small as the upper bound on generalization error is higher for smaller values of α.

Theorem 2.2. *Let $\hat{y} = \mathbf{f}_\theta(x)^T \psi$ be a linear regression problem as before. Let us further assume that $\mathbf{f}_\theta(x)\ \forall x \sim \mathcal{D}_{train}$ is a representative subset of the inputs from true data distribution: $(x, y) \sim \mathcal{D}$. Assuming a power-law distribution in the eigenspectrum of representations at $\mathbf{f}_\theta$, i.e. $\lambda_n = \frac{c}{n^\alpha} \quad \forall n \geq n^*$, where $n^* \in \{1, 2...N\}$, the generalization error after T weight update steps, $\mathcal{G}(T)$ is:*

$$\begin{aligned} \mathcal{G}(T) := \mathbb{E}_{x,y\sim\mathcal{D}}[(y - \mathbf{f}_\theta(x)^T\psi)^2] \quad &\leq \mathcal{O}\left(r_{\hat{d}}(\Sigma(\mathbf{f}_\theta)\right) \\ \textit{where} \qquad r_{\hat{d}}(\Sigma(\mathbf{f}_\theta)) = \frac{\sum_{i=\hat{d}}^m \lambda_i}{\sum_{i=1}^m \lambda_i} &= \frac{\sum_{i=\hat{d}}^m i^{-\alpha}}{\sum_{i=1}^m i^{-\alpha}} \end{aligned} \tag{2}$$

Here, $m = min(n, D)$ is the rank of $\Sigma_n(\mathbf{f}_\theta)$.

Notably, $r_{\hat{d}}(\Sigma(\mathbf{f}_\theta))$ can be thought of as a measure of effective rank of the covariance matrix and grows with decreasing values of α. Therefore, the upper bound for generalization error is higher for lower α. Taken together, Theorems 2.1 and 2.2 suggest that α can neither be too high nor too low and there exists a tradeoff region which results in efficient learning and strong generalizability, where these representations form a good basis for gradient-based optimization on downstream task performance.

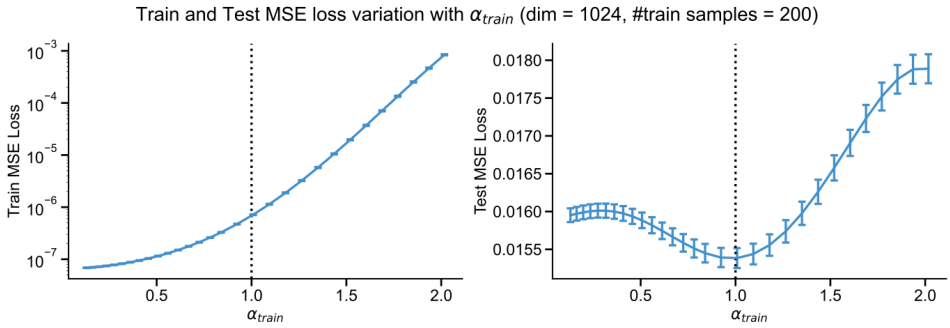

Figure 2: Overparameterized linear regression, with inputs drawn from Gaussian distribution with power-law in the covariance matrix. Models are trained with gradient descent with initialization $\psi_0 = \mathbf{0}$. Note that $\alpha > 1$ particularly suffers from high train, test MSE loss, with low generalization error with $\alpha \approx 1$.

Another motivating example Before exploring the link between α and generalization performance of deep networks, we validate our theorems in a simple linear regression setting.

We first consider its relationship to the finite dimensional regression setting as in [12]. Specifically, Theorem 6 of their paper states that the conditions for "benign overfitting", wherein a model can perfectly fit noisy training data without any subsequent loss of performance on testing data, may be looser in finite dimensions as opposed to the necessary and sufficient condition of $\alpha = 1$ in infinite dimensions. To empirically test this in the finite high dimensional setting, we examine linear least squares regression using different covariate structures for the input data. Formally, we consider covariates $\{x_i\}_{i=1}^N$, such that $x_i \in \mathbb{R}^d$ is sampled from a Gaussian distribution with covariance structure $\Sigma = \text{diag}\{\lambda_1, ...\lambda_d\}$ where $\lambda_j \sim PL(\alpha)$, i.e $\lambda_j \propto cj^{-\alpha}$. We assume access to the corresponding labels $\{y_i\}_{i=1}^N$ generated under a teacher function θ^*, such that $y_i = x_i^T\theta^* + \epsilon_i$. We find that in this scenario there is a clear relationship between the proximity of α to 1 and the presence of benign overfitting. As shown in Fig. 2, when α is close to 1, the training loss is low, but the validation loss is also low. Thus, when α is close to 1, the generalization properties are at their best. This example thus provides another hint that α is a potential measure for how well a model will be able to generalize.

3 Experimental Setup

Our theoretical results suggest the following: for representations with power-law characteristics, the decay-coefficient (α) effects the *adaptivity*, where the convergence rate of linear regression with gradient descent scales exponentially $\mathcal{O}(n^\alpha)$ with α. On the other hand, if trained sufficiently long, the upper bound for generalization error improves with larger α. Together, these items imply that if one wants to use gradient descent to train a downstream task then α for the representations used

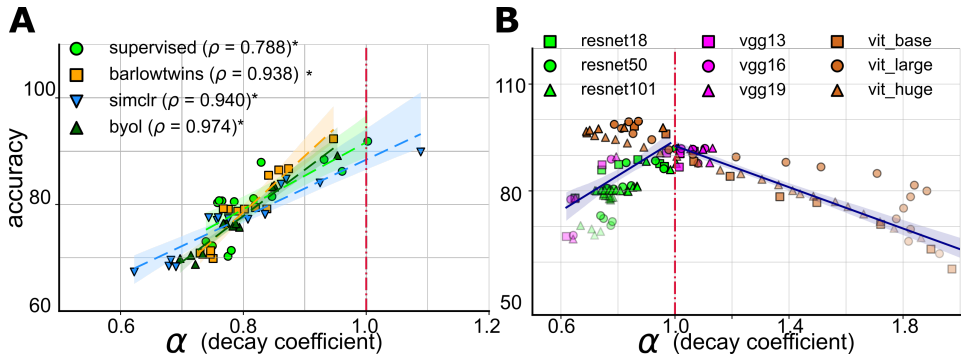

Figure 3: α **is predictive of out-of-distribution object recognition performance**. α is strongly correlated to object recognition performance on STL10 across different architectures and pretraining loss functions.

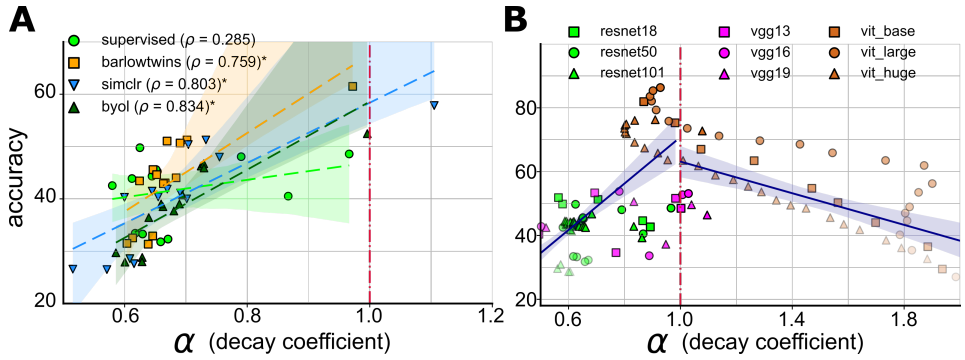

Figure 4: α **is predictive of out-of-distribution object recognition performance**. α is strongly correlated to object recognition performance on MIT67 across different architectures and pretraining loss functions.

should be neither too large nor too small, i.e. high quality representations are those with α in a "Goldilocks" zone. Building on these insights, we now design experiments to empirically study the potential for α to be used as a representation quality metric. We do so by measuring the relationship between α and the standard metric of representation quality, i.e. downstream task generalization performance. In this work, we restrict our scope to vision tasks, specifically object classification and scene recognition. In particular, our experiments:

- **measure the eigenspectrum decay & generalization** : We systematically evaluate *in-distribution* and *out-of-distribution* generalization, and simultaneously, we also measure α, across diverse choices for network architectures, pretraining learning objectives, and across multiple datasets. We then look at whether α values near some neighbourhood around 1 are indeed predictive of representations that generalize well to downstream tasks.
- **explore α as a metric for model-selection in SSL:** Measuring α is label-free and computationally inexpensive, as it requires a single forward-pass on the downstream dataset. We thus ask, is α predictive of generalization capabilities for current SSL algorithms, and if so, can it be used for model selection? We run extensive ablations on design of the BarlowTwins [1] algorithm in an effort to answer this question.

3.1 Measuring eigenspectrum Decay & Generalization

Evaluation Protocol: To measure α for the learned representation $\mathbf{f}(x_i)$, we first extract features from intermediate layers of a DNN pretrained on $\mathcal{D}_{pretrain}$ and estimate the corresponding covariance matrix $\Sigma_n(\mathbf{f})$. Next, we compute the full set of numerical eigenvalues, and estimate α by a fitting a linear model on the eigenspectrum in `log-log` scale. Our pretrained models are taken from PyTorch Hub [15] and `timm` [16]. Given that we observed no significant difference between the observed α values in the train and test sets, we refer to this empirical estimate as the α for the dataset.

To estimate the capacity of intermediate representations in solving the downstream task, we train a linear readout layer, $\mathbf{g}(.)$, from representations to target logits. Intuitively, this comes down to establishing a relationship between the manifold geometry and linear separability of representations [17]. Thereafter, we observe the correlation between estimated α and the linear readout performance. Notably, we also tried non-linear $\mathbf{g}(.)$ and observed a similar trend in results (see Appendix B.1).

Inductive bias of Model Architecture In this section, we investigate the relationship between α of the representation covariance matrix and object/scene recognition performance in DNNs with different backbones and the role of depth. In order to do so, we examine varying depth configurations within network architectures across three generations of models on the STL10 and MIT67 dataset [18]. We choose both object and scene recognition task as they fundamentally differ in the nature of the features that must be extracted to perform well. For scene recognition, the model focuses on global features in the entire image, while for object recognition, the model focuses on local features of the image containing the object[19].

We examine deep Convolutional Neural Networks (CNNs) without any residual connection as our first family of models. Specifically, we choose three different configurations of VGG-Net [20], namely VGG-13, VGG-16 and VGG-19. We inspect representations that are input to the dropout and MaxPool layers during the forward pass of the network. Second, we consider Deep Residual Networks [21] which are widely used in computer vision. We inspect the representations that are input to each of the residual blocks as well as the Adaptive average pool in ResNet-13, ResNet-50 and ResNet-101 during their respective forward passes. Finally, owing to the recent success of transformers in object recognition tasks we consider Vision Transformers (ViT) [22] as the third family of models, namely ViT-Base/8 , ViT-Large/16 and ViT-Huge/14. Unlike VGG and ResNet, we only look at features in the intermediate layers because it summarizes the entire input image and is used in practice for class prediction. For all model architectures, we use the weights obtained from pretraining on ImageNet.

Fig. 3 and Fig. 4 illustrate the relation between performance and α for all the nine architectures across intermediate layer representations, as described above. We found that, while most intermediate representations in CNNs (with or without residual connections) exhibit $\alpha < 1$, representations in ViTs mostly exhibit $\alpha > 1$. Nevertheless, representations extracted from the deepest layers of all the models exhibit α close to 1, irrespective of the total depth of each model (see Appendix B). Furthermore, the performance on downstream task increases with depth. This is unsurprising because all networks were trained to perform object recognition on ImageNet [23] and thereby would have leveraged hierarchical processing to learn features that are tuned towards object recognition. Surprisingly enough, we observe a strong significant correlation between α and performance on the STL10 dataset, i.e. a different data distribution than the training dataset, across layers and model architectures ($\rho = -0.922$, $^*p < 0.05$ for representations exhibiting $\alpha > 1$ and $\rho = -0.922$, $^*p < 0.05$ for representations exhibiting $\alpha < 1$). It is worth noting here that the correlation was weaker for the earliest layers of each model. We believe that early layers learn more task invariant features that reflect the statistics of natural images [24, 25] and therefore lack task relevant information in their representations. Taken together, this observation confirms our hypothesis that α is a good indicator of out-of-distribution generalization performance when representations possess task relevant information. Thus, α has a *necessary but not sufficient* relationship with generalization (Appendix C).

Learning objective In this section, we first aim to understand how the α value changes across the layers of a fixed architecture DNN when trained with different learning objectives. We take a ResNet-50 model [21] pre-trained using three different SSL algorithms, namely SimCLR [2], BYOL [26] and Barlow Twins [1], and the supervised learning loss objectives on ImageNet-1k[23] dataset. We use a similar procedure as before to extract representations from the network and estimate α.

Similar to results in the previous section, all networks irrespective of the pretraining loss function, exhibit α closer to 1 in the deeper layers in contrast to intermediate layers (see Appendix Fig. 9). This surprising result indicates that although the pre-training loss function was different, representations extracted from deepest layers are reflective of the object/scene semantics in natural images. Furthermore, Fig. 3 and Fig. 4 illustrate the strong correlation between α and generalization performance on STL10 and MIT67 across all pre-training loss functions. Together with results from the previous section, we validate our hypothesis that the representations that demonstrate good out-of-distribution generalization performance are characterized by α in the neighbourhood of 1. This suggests that high quality representations are indeed those with α in this region.

3.2 Label-agnostic metric for Model selection

With empirical evidence of α being a good measure for generalization, we now wish to study its suitability as a metric for identifying the best among models pretrained with SSL algorithms. A label-free metric like α could be beneficial when we don't have access to the downstream task annotations, and the SSL loss is not useful to distinguish models with good generalization performance (see Fig. 5). To investigate this, we exhaustively ablate the relationship between alpha and model performance across a wide range of hyper-parameters for a representative non-contrastive SSL algorithm.

Current SSL algorithms struggle with dimension collapse (characterized by large α), due to pathologies in training dynamics. To study α across a wide-range of values, for our model selection experiments we pick the *Barlow Twins*[1], a non-contrastive SSL algorithm that explicitly optimizes to avoid dimension collapse. In particular, the Barlow Twins learning objective ($\mathcal{L}_{\mathrm{BT}}$) proposes

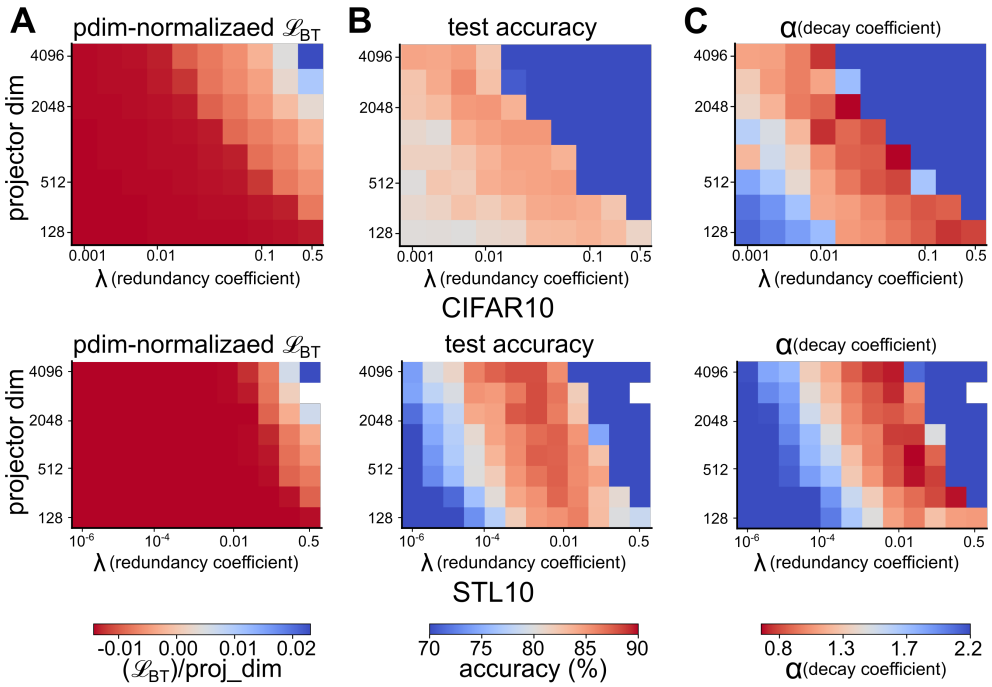

Figure 5: α **as a metric to inform model selection.** The SSL loss (training for same number of gradient steps) is no longer useful to distinguish models with superior downstream performance. However, decay coefficient α shows strong correspondence to downstream test accuracy over a large hyperparameter ranges. Measuring in-distribution generalization for (A-C) BarlowTwins trained and evaluated on CIFAR10. (D-F) BarlowTwins trained and evaluated on STL10.

imposing a *soft-whitening* constraint :

$$\mathcal{L}_{\text{BT}} = \underbrace{\sum_{i=1}^{d}(1-\mathcal{C}(\mathbf{f}_\theta)_{ii})^2}_{\text{invariance}} + \lambda \underbrace{\sum_{i=1}^{d}\sum_{j\neq i}^{d}\mathcal{C}(\mathbf{f}_\theta)_{ij}^2}_{\text{redundancy-reduction}} \quad \text{s.t. } \mathcal{C}(\mathbf{f}_\theta)_{ij} = \frac{\sum_n \mathbf{f}_\theta(x^A)_i \mathbf{f}_\theta(x^B)_j}{\sqrt{\sum_n \mathbf{f}_\theta(x^A)_i^2 \sum_n \mathbf{f}_\theta(x^B)_j^2}} \tag{3}$$

Notably with sufficient large λ, the model would impose an $\alpha = 0$ constraint on the representations [27]. Despite the intuitive connection between α and λ, it is unclear whether this relationship holds across a wide-range of values for model hyperparameters. To empirically establish this connection, we vary λ (redundancy coeffiecient), projection head dimensionality and learning rate, and train a ResNet50 encoder using *Barlow Twins* learning objective. We provide the results for CIFAR10 [28] and STL10 [29] in Fig. 5 demonstrating that α is a strong indicator of in-distribution generalization performance across a large range of hyperparameter ranges of $\mathcal{L}_{BT}$. We defer the reader to the Appendix Figs. 14 & 15 for similar trends in optimization hyperparameters, i.e. learning rate and weight decay parameters. Furthermore, α is predictive of out-of-distribution generalization performance in these settings (see Appendix B.3).

Model Selection on a Computer Budget: With a constrained compute budget, using α for model selection is significantly cheaper than evaluating the representations on a suite of downstream tasks. To illustrate this, consider the standard alternative of training a linear classifier on the representations and evaluating test accuracy. Compared to training a linear probe, which requires multiple (K)

Table 1: We report the compute time for CIFAR10, STL10, and ImageNet. While CIFAR10 and STL10 are trained for 200 epochs for downstream classification ImageNet is trained for 100 epochs. (tested in 1 A100)

Dataset	w/ eigenspectrum coefficient (α)	w/ linear probe	perf. gains
CIFAR10 (s)	**3**$\pm$ **0.2**	48$\pm$ 7	$\sim$ 16$\times$
STL10 (s)	**5**$\pm$ **0.5**	80$\pm$ 10	$\sim$ 16$\times$
ImageNet (mins)	**4**$\pm$ **0.8**	58$\pm$ 5	$\sim$ 15$\times$

epochs of forward & backward passes
through the training dataset's features,
to achieve reasonable estimates of downstream accuracy, computing α requires a single PCA step on the validation dataset's features. In Table 1, we provide estimated times for model-selection with α and downstream accuracy for CIFAR10, STL10 and ImageNet. A detailed algorithm is provided in Algorithm 1.

Algorithm 1 Model selection using α

```
# M: Number of models that can be trained in parallel
# H: Number of sequential steps of model training
# K: Number of top models to log in memory for model selection

α_min = 0
α_max = ∞
models_dict = {}

for iter in range(H):
   # train M models in parallel, each with different hyperparam config
   trained_models = train_SSL_models(num_parallel_models=M)

   # evaluate α for the trained models
   alpha_trained_models = evaluate_alpha(trained_models)
   best_trained_models = {model: alpha for model, alpha in
       alpha_trained_models if alpha ∈ [α_min,α_max]}
   models_stat = {model: run_linear_eval(model) for model, alpha in
       best_trained_models.items()}

   # Update α_min and α_max
   α_min = max({α_m | α_m > α_j & acc_m ≥ acc_j ∀ m,j ∈ models_stat})
   α_max = min({α_m | α_m < α_j & acc_m ≥ acc_j ∀ m,j ∈ models_stat})

   # Trim models_stat to keep only top K models
   threshold = get_best_models(model_statistics, topk=K)
   models_stat = {model: acc for model, acc in models_stat.items() if acc >
       threshold}

# Return best model performance from the 'selected' models
best_model_acc = max({acc_m ∀ m ∈ models_stat})
return best_model_acc
```

4 Related Work

Evaluating representations and model quality We note a substantial body of work aiming to empirically characterize structure of emergent representations in DNN, without requiring labels [30, 31]. One such index that quantifies the similarity of representations across layers (of same or different models) is Centered Kernel Alignment (CKA) [24]. While CKA does not provide explicit guidance for downstream performance, [32] show that in-distribution generalization gap can be predicted using a different index based on the model's parameters. In particular, they show that Empirical Spectral Density (ESD) of weight matrices for many DNNs obey a power-law, with the coefficient of decay being predictive of in-distribution performance. In the present work, we explore similar indices that potentially correlate with out-of-distribution generalization by examining the eigenspectrum of *activations*.

Generalization in Overparameterized Models Modern neural networks often have significantly more parameters than number of training samples, challenging classical understanding of the bias-variance tradeoff. Overparameterization permits neural networks to overfit to noise in training data, without impairing their generalization to unseen data. In recent work, Bubeck *et al.* proved that for the interpolator (function implemented by the neural network) to be smooth, overparameterization is a necessary condition [33].

Furthermore, this *benign overfitting* phenomenon in an overparameterized linear regression problem has been linked to the power law coefficient of the input covariance matrix [12]. Specifically, Bartlett *et al.* showed that for an infinite-dimensional linear regression problem, benign overfitting is possible iff the eigenspectrum satisfies a power law (upto polylog factors). More recently, Lee *et al.* found that the tail eigenvalues of infinite-width network kernels exhibit a power law decay [34]. Following this, Tripuraneni *et al.* explored high-dimensional random feature regression settings and

analytically showed a dependence between eigenspectrum decay rate of the feature covariance matrix and generalization error [35]. While these characterizations provide a theoretical understanding of generalization error in the asymptotic or random feature settings, corresponding questions in the finite dimensional DNN trained with gradient descent are open problems.

For deep linear networks trained using gradient descent, the eigenvalues of input covariance determine the generalization error dynamics [13]. Advani *et al.* demonstrated that small eigenvalues determine the convergence of training dynamics as well as the overfitting error at convergence. For overparameterized 2-layer neural networks, Arora *et al.* provided a fine grained analysis of generalization bounds [36]. In contrast, we study modern DNN architectures and explore the covariance structure of their learned features on visual recognition tasks.

5 Discussion

Summary Our experiments suggest a strong correlation between decay coefficient for sample eigenspectrum of representations, α, and both in-distribution, out-of-distribution generalization performance on tasks central to computer vision.

Representation Quality in High-Dimensional SSL We demonstrate that assessing the quality of a pretrained model with our label-free metric (α) is consistent with downstream generalization. While being computationally efficient, our procedure removes the dependence on labels, making the model selection more robust and amenable for privacy critical applications.

Necessary, but not sufficient condition : Notably, a task-agnostic measure like α is a necessary but not a sufficient condition for assessing good performance. To elaborate on this point, let us present a thought experiment. Suppose that we have our ideal representation space that satisfies our claims of exhibiting alpha close to the Goldilocks zone. If we perform a set of permutation operations on individual datapoint representations, the structure of the representation space remains the same but the mapping from representation to label is now destroyed. In doing this set of permutations, we have now arrived at a different set of representations that would exhibit the same (or similar) alpha but demonstrate a lower task performance when measured using a linear readout layer. A similar argument is presented as theorem C.1 in [37]. Extending the conclusions of this thought experiment to our observations, it is clear that we could have models with similar alpha values but distinctly different performance. But an alpha value that is not in the Goldilocks zone would be associated with inferior model performance. This property is the core of our claim and allows us to propose alpha as a measure for model selection in SSL pipelines.

Redundancy in ViT representations Unlike other models that we considered, representations from early layers of ViT had a rapid eigenspectrum decay with $\alpha >> 1$ (see Fig. 3 and Fig. 4). The transformer architecture has a notable difference by design, i.e. early layers possess global receptive field context via self-attention on patch embeddings. Raghu *et al.* found that early ViT layers incorporate both local and global information [31]. Based on these insights, one intuitive interpretation of our results is that the representations have a low effective rank, and encode redundant information relevant across multiple scales.

Limitations While the role of α and its relationship to generalization performance is better understood in the asymptotic setting for linear regression, similar questions in finite dimensional nonlinear models are unanswered. It is also worth noting that the empirical correlation was weaker for the earliest layers in each model. We believe that early layers learn more task invariant features such as corners and edges [24, 25], and thereby lack rich semantic information for fine-grained downstream task. In this case, distinguishing between poorly-trained models may be inconsistent with α.

Future Directions Learning efficiently at scale from unlabelled datasets poses an exciting open problem in deep representation learning. We hope this work encourages new perspectives into model selection for SSL pipelines, and informs design of learning objectives and model architectures to learn *task-agnostic, adaptive* features. Understanding the behaviour of α (notably, a scale-invariant metric) in high-dimensional representation learning might provide insight towards developing a theoretically grounded understanding of generalization in deep neural networks.

Acknowledgements

The authors would like to thank Zahraa Chorghay and Colleen Gillon for their aesthetic contribution to the manuscript and figures. This research was enabled in part by support provided by Mila and Compute Canada. This work was supported by Vanier Canada Graduate scholarship (AG); Healthy Brains, Healthy Lives (AG & BAR); NSERC, Grant No. RGPIN-2020-05105 and RGPAS-2020-00031 (BAR); and CIFAR, Canada AI Chair & Learning in Machines and Brains Fellowship (BAR).

References

[1] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021.

[2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.

[3] Yuandong Tian, Lantao Yu, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning with dual deep networks. *arXiv preprint arXiv:2010.00578*, 2020.

[4] Daniel J Felleman and David C Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1):1–47, 1991.

[5] Claus C Hilgetag and Alexandros Goulas. 'hierarchy'in the organization of brain networks. *Philosophical Transactions of the Royal Society B*, 375(1796):20190319, 2020.

[6] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D Harris. High-dimensional geometry of population responses in visual cortex. *Nature*, 571(7765):361–365, 2019.

[7] Weihao Kong and Gregory Valiant. Spectrum estimation from samples. *The Annals of Statistics*, 45(5):2218–2247, 2017.

[8] Nathan CL Kong, Eshed Margalit, Justin L Gardner, and Anthony M Norcia. Increasing neural network robustness improves match to macaque v1 eigenspectrum, spatial frequency preference and predictivity. *PLOS Computational Biology*, 18(1):e1009739, 2022.

[9] J Nassar, P Sokol, S Chang, and K Harris. On 1/n neural representation and robustness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

[10] Zijian Jiang, Jianwen Zhou, and Haiping Huang. Relationship between manifold smoothness and adversarial vulnerability in deep learning with local errors. *Chinese Physics B*, 30(4):048702, 2021.

[11] Zeke Xie, Qian-Yuan Tang, Yunfeng Cai, Mingming Sun, and Ping Li. On the power-law spectrum in deep learning: A bridge to protein science. *arXiv preprint arXiv:2201.13011*, 2022.

[12] Peter L Bartlett, Philip M Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.

[13] Madhu S Advani, Andrew M Saxe, and Haim Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428–446, 2020.

[14] Vatsal Shah, Anastasios Kyrillidis, and Sujay Sanghavi. Minimum norm solutions do not always generalize well for over-parameterized problems. *stat*, 1050:16, 2018.

[15] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,

Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[16] Ross Wightman. Pytorch image models. https://github.com/rwightman/pytorch-image-models, 2019.

[17] SueYeon Chung, Daniel D Lee, and Haim Sompolinsky. Classification and geometry of general perceptual manifolds. *Physical Review X*, 8(3):031003, 2018.

[18] Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 413–420, 2009.

[19] A Oliva and A Torralba. Chapter 2 building the gist of a scene: the role of global image features in recognition. *Progress in Brain Research*, pages 23–36, 2006.

[20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[21] 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*, pages 770–778, 2016.

[22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[24] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pages 3519–3529. PMLR, 2019.

[25] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.

[26] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020.

[27] Bobby He and Mete Ozay. Exploring the gap between collapsed & whitened features in self-supervised learning. In *International Conference on Machine Learning*, pages 8613–8634. PMLR, 2022.

[28] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Master's thesis*, 2009.

[29] 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*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.

[30] Thao Nguyen, Maithra Raghu, and Simon Kornblith. Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. *arXiv preprint arXiv:2010.15327*, 2020.

[31] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34, 2021.

[32] Charles H Martin, Tongsu Serena Peng, and Michael W Mahoney. Predicting trends in the quality of state-of-the-art neural networks without access to training or testing data. *Nature Communications*, 12(1):1–13, 2021.

[33] Sébastien Bubeck and Mark Sellke. A universal law of robustness via isoperimetry. *arXiv preprint arXiv:2105.12806*, 2021.

[34] Jaehoon Lee, Samuel Schoenholz, Jeffrey Pennington, Ben Adlam, Lechao Xiao, Roman Novak, and Jascha Sohl-Dickstein. Finite versus infinite neural networks: an empirical study. *Advances in Neural Information Processing Systems*, 33:15156–15172, 2020.

[35] Nilesh Tripuraneni, Ben Adlam, and Jeffrey Pennington. Covariate shift in high-dimensional random feature regression. *arXiv preprint arXiv:2111.08234*, 2021.

[36] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.

[37] Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 19250–19286. PMLR, 17–23 Jul 2022.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See the Discussion section (Section 5).
   (c) Did you discuss any potential negative societal impacts of your work? [No] Our work proposes a new measure of evaluating representation quality, specifically representations that have the potential of training quickly and generalizing well to downstream tasks. However, in our current efforts we do not attempt to address growing concerns about bias in deep network representations. We believe that our proposal could provide a platform for future work to investigate the utility of α for inspecting bias in representations. Taken together, our work could form a path to alleviate some of the potential negative impact. We leave this to future scope.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] We present detailed proofs along with a discussion of how our theoretical results are tied to previous theoretical advances in the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We provide implementation details for each experiment and generating each figure in the relevant section in Appendix.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] We have cited the assets that we used in the experiments.

(b) Did you mention the license of the assets? [No] The license is publicly available on each assets GitHub page or website.

(c) Did you include any new assets either in the supplemental material or as a URL? [No] We did not produce any new assets.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] The assets we used are open-source or openly available.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] The data that we are using clearly does not contain such content.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]