

489 **A Optimal Embeddings**

490 Recall the setting of self supervised learning as described in Balestriero and LeCun (5): given  
 491 a dataset  $\mathbf{X}' = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N \times D'}$  we construct a new dataset by creating  $k$  randomly  
 492 augmented views of the original data,  $\mathbf{X} = [\text{view}_1(\mathbf{X}'), \dots, \text{view}_k(\mathbf{X}')] \in \mathbb{R}^{Nk \times D}$ . The advantage  
 493 of doing so is that we can now leverage the knowledge that different views of the same underlying  
 494 datapoint are *semantically related*. We can express this notion of similarity in the symmetric matrix  
 495  $\mathbf{G} \in \{0, 1\}^{Nk \times Nk}$  with  $G_{ij} = 1$  if augmented datapoints  $i$  and  $j$  are semantically related (and  
 496  $G_{ii} = 1$  as any datapoint is related to itself). We can normalize  $\mathbf{G}$  such that its rows and columns  
 497 sum to 1 (so rows of  $\mathbf{G}$  are  $k$ -sparse with nonzero entries equal to  $1/k$ ).

498 Now let  $\mathbf{Z} \in \mathbb{R}^{Nk \times d}$  be an embedding of the augmented dataset. Then we have  $\mathbf{GZ} = [\mathbf{C}, \dots, \mathbf{C}]^T$   
 499 where  $\mathbf{C}$  is the matrix of centroid vectors introduced above, and the number of repetitions of  $\mathbf{C}$  is  $k$ .  
 500 Then because  $\sigma([\mathbf{C}, \dots, \mathbf{C}]) = \sqrt{k}\sigma(\mathbf{C})$  we can write MMCR loss function as,

$$\begin{aligned} \mathcal{L} &= -\|\mathbf{GZ}\|_* \\ &= -\|\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T\mathbf{U}\mathbf{S}\mathbf{V}^T\|_* \\ &= -\|\mathbf{\Lambda}\mathbf{Q}^T\mathbf{U}\mathbf{S}\|_* \end{aligned} \tag{6}$$

501 Where we have taken the eigendecomposition of  $\mathbf{G}$  which is real and symmetric and the SVD of  
 502  $\mathbf{Z}$ , and then used the fact that the singular value spectrum is invariant under left or right orthogonal  
 503 transformations. We now show that a global optima of this objective is achieved when the left singular  
 504 vectors of  $\mathbf{Z}$  are the eigenvectors of  $\mathbf{G}$  and the singular values of  $\mathbf{Z}$  are proportional to the eigenvalues  
 505 of  $\mathbf{G}$ . Throughout we will assume that the size of the dataset is greater than the dimensionality of the  
 506 embeddings,  $N > d$ , as is the case in practical applications. First we prove a simple lemma about  
 507 the spectrum of matrices who are extended by zeros (i.e. embedded in a higher dimensional space).

508 **Lemma A.1:** For  $\mathbf{A} \in \mathbb{R}^{N \times N}$ ,  $\mathbf{B} \in \mathbb{R}^{N \times d}$  with  $d < N$ ,  $\|\mathbf{AB}\|_* = \|\mathbf{A}\tilde{\mathbf{B}}\|_*$  where  $\tilde{\mathbf{B}} = [\mathbf{B}, \mathbf{0}] \in$   
 509  $\mathbb{R}^{N \times N}$ .

510 **Proof:** First note that  $\mathbf{A}\tilde{\mathbf{B}} = [\mathbf{AB}, \mathbf{0}]$  so it suffices to show that for arbitrary  $\mathbf{X}$  that  $\sigma(\mathbf{X}) =$   
 511  $\sigma([\mathbf{X}, \mathbf{0}])$ . Taking the SVD of  $\mathbf{X}$ ,

$$\mathbf{X} = \begin{bmatrix} \mathbf{U} & \tilde{\mathbf{U}} \end{bmatrix} \begin{bmatrix} \mathbf{\Sigma} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{V}^T \end{bmatrix} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$$

512 Then a valid singular value decomposition for  $\tilde{\mathbf{X}}$  is

$$\tilde{\mathbf{X}} = \begin{bmatrix} \mathbf{U} & \tilde{\mathbf{U}} \end{bmatrix} \begin{bmatrix} \mathbf{\Sigma} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{V}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$$

513 Clearly then,  $\|\mathbf{X}\|_* = \|\tilde{\mathbf{X}}\|_*$

514 **Theorem:** The proposed loss achieves a global minimum when the left singular vectors of  $\mathbf{Z}$  are the  
 515 eigenvectors of  $\mathbf{G}$ , and the singular values of  $\mathbf{Z}$  are proportional to the top  $d$  eigenvalues of  $\mathbf{G}$ .

516 **Proof:** Let  $\tilde{\mathbf{Z}} = [\mathbf{Z}, \mathbf{0}] \in \mathbb{R}^{Nk \times N}$ . By Lemma A.1 we have  $\|\mathbf{GZ}\|_* = \|\mathbf{G}\tilde{\mathbf{Z}}\|_*$ . Von Neumann's  
 517 trace inequality can be used to show  $\|\mathbf{G}\tilde{\mathbf{Z}}\|_* \leq \sum_{i=1}^{Nk} \sigma_i(\mathbf{G})\sigma_i(\tilde{\mathbf{Z}})$  (see Marshall et al. (48) for  
 518 proof). Examining (4) it is clear that this bound is achieved when  $\mathbf{U} = \mathbf{Q}$ . The problem can therefore  
 519 be reduced to the constrained optimization problem,

$$\begin{aligned} &\min_{\sigma_i(\tilde{\mathbf{Z}})} \sum_{i=1}^{Nk} \sigma_i(\mathbf{G})\sigma_i(\tilde{\mathbf{Z}}) \\ &\text{subject to } \sum_{i=1}^{Nk} \sigma_i(\tilde{\mathbf{Z}})^2 = Nk \end{aligned}$$

520 where the constraint comes from the fact that columns of  $\mathbf{Z}$  are unit vectors. Intuitively, we are  
 521 maximizing the inner product between a fixed vector  $\sigma(\mathbf{G})$  and a vector with fixed L2 norm. The

522 solution of course is to align the two vectors as closely as possible, i.e. when  $\sigma_i(\tilde{\mathbf{Z}}) \propto \sigma_i(\mathbf{G})$  for  
 523  $i = 1, \dots, d$ . It is worth noting that by construction  $\sigma_i(\tilde{\mathbf{Z}}) = 0$  for  $i > d$  and the columns of  $\mathbf{U}$   
 524 associated with these zero valued singular values are unconstrained.

## 525 B Pytorch Style Pseudocode for MMCR

```

526 1 # h: encoder
527 2 # g: projection head
528 3 # T: momentum temperature
529 4 # B: batch size
530 5 # K: number of augmentations
531 6 # D: projector output dimensionality
532 7 #
533 8 # lambda: trade-off parameter
534 9
535 0 f_o, g_o = ResNet50(), MLP() # online networks
536 1
537 2 # initialize momentum network with identical params
538 3 f_m, g_m = f_o.copy(), g_o.copy()
539 4
540 5 # momentum networks are not updated via gradient descent
541 6 f_m.requires_grad = False
542 7 g_m.requires_grad = False
543 8
544 9 for x in loader:
545 0     # K randomly augmented views
546 1     x = multi_augment(x) # B x K x H x W
547 2
548 3     # push through encoder and projector
549 4     z_o = g_o(h_o(x)) # B x K x D
550 5     z_m = g_m(h_m(x)) # B x K x D
551 6     z = concatenate(z_o, z_m, dim=1) # append outputs
552 7
553 8     # project onto unit sphere
554 9     z = normalize(z, dim=-1)
555 0
556 1     # calculate centroids (mean over augmentation axis)
557 2     c = z.mean(dim=1) # B x D
558 3
559 4     # calculate singular values
560 5     U_z, S_z, V_z = svd(z) # batch svd
561 6     U_c, S_c, V_c = svd(c)
562 7
563 8     # calculate loss
564 9     loss = -1.0 * sum(S_c) + lambda * sum(S_z) / B
565 0
566 1     # backward pass and optimization step
567 2     loss.backward()
568 3     optim.step()
569 4
570 5     # perform momentum update
571 6     with torch.no_grad():
572 7         f_m.parameters() = (1 - T) * f_o.parameters() + T * f_m.
573         parameters()
574 8         g_m.parameters() = (1 - T) * g_o.parameters() + T * g_m.
575         parameters()

```

## 576 C Mean Field Theory Manifold Capacity Background Information

577 For completeness we summarize some of the central arguments from Chung et al. (16), which  
 578 develops the general form of manifold capacity theory.

579 **Mean Field Theory** Recall the problem setting for manifold capacity analysis: given a set of  $P$   
580 manifolds embedded in a feature space of dimensionality  $D$ , each assigned a random binary class  
581 label (16). Manifold capacity theory is concerned with the question: what is the largest value of  
582  $\frac{P}{D}$  such that there exists (with high probability) a hyperplane separating the two classes? In the  
583 thermodynamic limit, where  $P, D \rightarrow \infty$  but  $\frac{P}{D}$  remains finite, the inverse capacity can be written  
584 exactly,

$$\alpha_M^{-1} = \mathbb{E}_{\vec{T}}[F(\vec{T})] \quad (7)$$

585 where,  $F(\vec{T}) = \min_{\vec{V}} \left\{ \|\vec{V} - \vec{T}\|^2 \mid g_{\mathcal{S}}(\vec{V}) \geq 0 \right\}$ ,  $\mathcal{S}$  is the set defining the manifold geometry (i.e.  
586 the set of vectors  $\vec{S}$  that are points on an individual manifold),  $\vec{T}$  are random vectors drawn from a  
587 white multivariate Gaussian distribution, and  $g_{\mathcal{S}}(\vec{V}) = \min_{\vec{S} \in \mathcal{S}} \{ \vec{V} \cdot \vec{S} \}$ , is the concave support  
588 function.

589 The KKT equations for this convex optimization problem are:

$$\begin{aligned} \vec{V} - \vec{T} - \lambda \tilde{S}(\vec{T}) &= 0 \\ \lambda &\geq 0 \\ g_{\mathcal{S}}(\vec{V}) - \kappa &\geq 0 \\ \lambda [g_{\mathcal{S}}(\vec{V}) - \kappa] &= 0. \end{aligned} \quad (8)$$

590 , where  $\tilde{S}(\vec{T})$  is a subgradient of the support function. When the support function is differentiable,  
591 the subgradient is unique and equal to the gradient,

$$\tilde{S}(\vec{T}) = \nabla g_{\mathcal{S}}(\vec{V}) = \arg \min_{\vec{S} \in \mathcal{S}} \vec{V} \cdot \vec{S} \quad (9)$$

592  $\tilde{S}(\vec{T})$  is the unique point in the convex hull of  $\mathcal{S}$  that satisfies the first KKT equation, and is called  
593 the ‘‘anchor point’’ for  $\mathcal{S}$  induced by the random vector  $\vec{T}$ .

594 **Equivalent Interpretation of Anchor Points** For a given dichotomy (random binary class labelling)  
595 the weight vector of the maximum margin separating hyperplane can be decomposed into a sum of at  
596 most  $P$  vectors, with each manifold contributing a single vector, which lies within the convex hull of  
597 the manifold. The position of said point is a function of the manifolds position relative to all  
598 of the other manifolds in the space and depends on the particular set of random labels. Thus there  
599 exists a distribution of separating-hyperplane-determining-points for each individual manifold. Using  
600 the cavity method it can be shown that these points are none other than the anchor points that are  
601 involved in solving the optimization problem described above (31).

602 **Numerical Solution** To solve the mean field equations numerically, one samples several random  
603 Gaussian vectors  $\vec{T}$ , and then for each  $\vec{T}$ ,  $\vec{V}$  and  $\vec{S}$  are determined by solving the quadratic program-  
604 ming program given above. The capacity is then estimated as the mean value of  $F$  or the samples  
605  $\vec{T}$ .

606 **Manifold Geometries** The way the capacity varies in terms of the statistics of the anchor points can  
607 be simplified by introducing two key quantities, the manifold radius  $R_M$  and manifold dimensionality  
608  $D_M$ :

$$\begin{aligned} R_M^2 &= \mathbb{E}_{\vec{T}}[\|\tilde{S}(\vec{T})\|^2] \\ D_M &= \mathbb{E}_{\vec{T}}[\vec{T} \cdot \hat{S}(\vec{T})] \end{aligned} \quad (10)$$

609 where  $\hat{S}(\vec{T})$  is a unit-vector in the direction of the anchor point  $\tilde{S}$ . In particular as discussed in the  
610 main text, the manifold capacity can be approximated by  $\phi(R_M \sqrt{D_M})$  where  $\phi$  is a monotonically  
611 decreasing function.

612 **Elliptical Geometries** In the case where the manifolds exhibit elliptical symmetries, the manifold  
 613 radius and dimensionality can be written in terms of the eigenvalues of the covariance matrix of the  
 614 anchor points:

$$R_M^2 = \sum_i \lambda_i^2$$

$$D_M = \frac{(\sum_i \lambda_i)^2}{\sum_i \lambda_i^2} \tag{11}$$

615 So, in this case  $R_M$  is the total variability of the anchor points, and  $D_M$  is a generalized participation  
 616 ratio of the anchor point covariance, a well known soft measure of dimensionality.

## 617 D Additional Pre-training information

618 **Settings for CIFAR/STL-10** We take the parameters of each augmentation directly from Zbontar  
 619 et al. (63), but for these lower resolution images we omitted Gaussian blurring and solarization  
 620 augmentations. All models were trained for 500 epochs using the Adam optimizer (40) with a  
 621 learning rate of  $1e - 3$  and weight decay of  $1e - 6$ . For all three methods we used a one hidden layer  
 622 MLP with hidden dimension of 512 and output dimension of 128 for the projector head  $g$ . We swept  
 623 batch size for each method and chose the one that resulted in the highest downstream task performance.  
 624 For both SimCLR and Barlow Twins we found that a batch size of 128 was optimal (among 32, 64,  
 625 128, 256, and 512) for all 3 datasets. For MMCR there is a trade-off between batch size and the  
 626 number of augmentations used, and the optimal value of that trade-off is highly dataset dependent. For  
 627 CIFAR-10 and CIFAR-100 we used batch size of 32 and 40 views, and for STL-10 we used a batch of  
 628 64 with 20 views For Barlow Twins we used  $\lambda = \frac{1}{128}$  which normalizes for the number of elements  
 629 in the on-diagonal and off-diagonal terms in the loss. For SimCLR we used the recommended  
 630 setting of  $\tau = 0.5$ . The overall performance of both baseline methods (and likely MMCR as well)  
 631 could be increased with a more thorough hyperparameter search and by employing methodology that  
 632 more closely matches the original works. For example, both methods would likely benefit from the  
 633 combination of larger batch size, the use of the LARS optimizer (which is designed for large batch  
 634 optimization), a learning rate scheduler consisting of linear warm-up followed by cosine annealing,  
 635 longer training, and the use of more diverse augmentations (i.e. including solarization and gaussian  
 636 blur). Additionally Barlow Twins reports that the representation can benefit from using a much larger  
 637 projector network than we use. Because our goal was primarily to demonstrate that MMCR can  
 638 produce representations that are comparable to these baselines rather than to produce state-of-the-art  
 639 results on small scale datasets we opted for simplifications wherever possible (using off the shelf  
 640 Adam for optimization with a fixed learning rate, and fixing architectural hyperparameters like the  
 641 projector dimensionality).

642 **Settings for ImageNet-100** For ImageNet we more closely match the pre-training procedures of  
 643 previous works. We use a batch size of 2048 and a smaller number of views for MMCR (4), and  
 644 also use the full suite of augmentations from Zbontar et al. (63). For the sake of efficiency we train  
 645 for a reduced number of epochs (200). For MMCR and SimCLR we modified the projector hidden  
 646 dimensionality to be 4096 for the projector head, following the original work (12). For Barlow  
 647 Twins we used the recommended 2-layer MLP with hidden and output dimensions of 8192, and set  
 648  $\lambda = 5e - 3$ , however these hyperparameters were optimal for the full ImageNet dataset, and not  
 649 necessarily for ImageNet-100. We were unable to achieve better downstream performance using a  
 650 ResNet-50 backbone than what has previously been reported in the literature for this dataset with a  
 651 ResNet-18 backbone, therefore we report the ResNet-18 performance reported in (20). For SimCLR  
 652 we use  $\tau = 0.1$  which is the recommended setting for larger batch sizes.

653 **Settings for ImageNet-1k:** For ImageNet-1k we use mostly identical settings to ImageNet-100,  
 654 but we increased the capacity of the projector network (using a 2 hidden layer MLP with hidden  
 655 dimenisons of 8192 and output dimension of 512). We scaled the learning rate linearly with batch  
 656 size:  $lr = 0.6 \times \frac{\text{batch size}}{256}$ . Additionally we reduce the number of pretraining epochs to 100. Finally  
 657 for ImageNet-1k we found that employing a momentum encoder slightly boosted downstream per-  
 658 formance (around +0.5% on ImageNet frozen-linear evaluation). Specifically, an identical encoding  
 659 network and projector architecture is initialized with the same parameters as the initial "online"  
 660 network, and during training the weights of this "momentum" network track a slowly moving average  
 661 of the online network parameters (only the online network parameters are updated via gradient

662 descent). Each augmented view is passed through both the online and momentum networks, and the  
663 resultant embeddings are all averaged to form the centroid vector for a particular image in the batch.  
664 We used a momentum coefficient of 0.99.

665 Pre-training on 16 A100 GPUs using 8 views (our most compute intensive setting) takes approximately  
666 32 hours.

## 667 E Details of Representational Analyses

### 668 E.1 Manifold Capacity Analysis

669 For each pre-trained model, we extract layer activations across the ResNet hierarchy after a forward  
670 pass of a set of images. For class manifold analysis, the set of images contain 10 classes, where each  
671 class has 100 examples. Augmentation manifolds instead have 100 exemplars with 100 examples  
672 each. Following (18), we take activations from all convolutional layers in ResNet-50 after a ReLU  
673 non-linearity. The specific extracted layers highlighted in bold fonts are given by Table 3. The  
674 final analysis results are averaged over five data samplings with different random seeds and random  
675 projections of intermediate features to lower-dimension spaces (default 5000 dimensions).

### 676 E.2 Gradient Coherence Analysis

677 In Fig. 3, for each of the classes of CIFAR-10, we generate 100 batches of 32 augmentation manifolds  
678 of samples from a specific class (with 40 augmentations each). We then measure the gradient of the  
679 loss function for each batch during different stages of training, and compute the cosine similarity  
680 between every pair of gradients. Across all stages of training the mean cosine similarity between  
681 gradients generated from batches of the same class is larger than those from distinct classes (left  
682 column). This observation remains true when isolating the gradients of parameters from different  
683 stages of in the resnet-50 hierarchy (center and right columns, respectively).

### 684 E.3 Manifold Subspace Alignment

685 For Fig. 4 we generated 100 samples from the augmentation manifolds of 500 images in the CIFAR-10  
686 dataset. We then measure the mean subspace angle (left column), fraction of shared variance (middle  
687 column) and centroid cosine similarity between each pair of manifolds. The same procedure was  
688 used for generating the data for Fig. 5.

689 **Subspace Angle.** Besides measuring the size and dimensionality of individual object manifolds we  
690 also wish to characterize the degree of overlap between pairs of manifolds. For this, we measure the  
691 angle between their subspaces (41), which is a generalization of the notion of angles that applies to  
692 subspaces of arbitrary dimension.

693 **Shared Variance.** Object manifolds will generally have a lower intrinsic dimensionality than the  
694 space in which they are embedded. Therefore, the data will have low variance along several of the  
695 principal vectors used to calculate the set of subspace angles, and so many of the principal angles  
696 will have little meaning. To address this limitation we also compute the shared variance between the  
697 linear subspaces that contain object manifolds.

## 698 F Implicit MMCR Effectively Reduces Augmentation Manifold Nuclear 699 Norm

700 To test whether or not implicit manifold compression actually reduces the mean augmentation  
701 manifold nuclear norm, we can vary the value of  $\lambda$ . Below we see the evolution of both terms of the  
702 loss for several different values of lambda during training on CIFAR-10. For these experiments the  
703 batch size was 64 and the number of augmentations per image was 4.0. As shown in Fig. 6, the level  
704 of compression of individual manifolds is nearly the same across all values of the parameter.

Table 3: A Total of 18 Extracted ResNet-50 Layers (in **Bold**) for MFTMA Analysis

Layer	Type	Conv2d Size (H × W × C)
pixel	<b>Input</b>	None
conv1	$\begin{bmatrix} [L] \text{Conv2d} \\ \text{BatchNorm} \\ \mathbf{ReLU} \end{bmatrix} \times 1$	$[[L] 7 \times 7 \times 64] \times 1$
conv2_x	$\begin{bmatrix} [L] \begin{bmatrix} [L] \text{Conv2d} \\ \text{BatchNorm} \\ \mathbf{ReLU} \end{bmatrix} \times 3 \end{bmatrix} \times 3$	$\begin{bmatrix} [L] 1 \times 1 \times 64 \\ 3 \times 3 \times 64 \\ 1 \times 1 \times 256 \end{bmatrix} \times 3$
conv3_x	$\begin{bmatrix} [L] \begin{bmatrix} [L] \text{Conv2d} \\ \text{BatchNorm} \\ \mathbf{ReLU} \end{bmatrix} \times 3 \end{bmatrix} \times 4$	$\begin{bmatrix} [L] 1 \times 1 \times 128 \\ 3 \times 3 \times 128 \\ 1 \times 1 \times 512 \end{bmatrix} \times 4$
conv4_x	$\begin{bmatrix} [L] \begin{bmatrix} [L] \text{Conv2d} \\ \text{BatchNorm} \\ \mathbf{ReLU} \end{bmatrix} \times 3 \end{bmatrix} \times 6$	$\begin{bmatrix} [L] 1 \times 1 \times 256 \\ 3 \times 3 \times 256 \\ 1 \times 1 \times 1024 \end{bmatrix} \times 6$
conv5_x	$\begin{bmatrix} [L] \begin{bmatrix} [L] \text{Conv2d} \\ \text{BatchNorm} \\ \mathbf{ReLU} \end{bmatrix} \times 3 \end{bmatrix} \times 3$	$\begin{bmatrix} [L] 1 \times 1 \times 512 \\ 3 \times 3 \times 512 \\ 1 \times 1 \times 2048 \end{bmatrix} \times 3$




Figure 6: Validation loss values for different values of  $\lambda$

## 705 G Classification Evaluation Procedure

706 **CIFAR and STL-10:** During pre-training all models were monitored with a k-nearest neighbor  
707 classifier (k=200) and checkpointed every 5 epochs. After pre-training, we trained linear classifiers  
708 on all checkpoints whose monitor accuracy was within 1% of the highest observed accuracy, and  
709 select the model that achieves the highest linear classification accuracy. Linear classifiers were trained  
710 using the Adam optimizer with batch size of 1024 and an initial learning rate of 0.1, which decayed  
711 according to a cosine scheduler over the course of 50 epochs. For the linear classifier training, at train  
712 time we use the same set of augmentations as during unsupervised pretraining, at test time we only  
713 use center cropping and random horizontal flipping.

714 **ImageNet-1k/100:** For ImageNet datasets we closely followed the most widely adopted evaluation  
 715 procedure. Following pre-training we freeze the encoder weights and train a linear layer in a  
 716 supervised fashion using SGD with a batch size of 256, learning rate of 0.3, and weight decay of 1e-6  
 717 for 100 epochs. During linear classifier training the only data augmentations are random cropping  
 718 and random horizontal flips, and during evaluation inputs are center cropped.

719 **Semi-Supervised:** For semi-supervised evaluation we mostly follow the procedure outlined in Bardes  
 720 et al. (6). We use the SGD optimizer with momentum of 0.9 and weight decay of 1e-6 and the  
 721 standard cross entropy loss. The augmentation procedure was the same as described above. Because  
 722 the linear classifier is being trained from scratch and the representation is being fine tuned the learning  
 723 rate for the parameters of the backbone is scaled down by a factor of 10, and both learning rates  
 724 (backbone and classifier) followed a cosine decay schedule for 20 epochs. We used a batch size of  
 725 256 and swept the learning rate over [0.1, 0.3, 1.0] for each model.

726 **Other Downstream Classification Tasks:** Classifiers on these datasets were trained in a similar  
 727 fashion to those trained on the CIFAR and STL-10 datasets. The only difference was that we trained  
 728 for only 20 epochs (which we found sufficient for convergence), the initial learning rate over [3e-2,  
 729 3e-3, 3e-4] for each model, and the augmentation procedure matched the standard setup for ImageNet  
 730 training (only random cropping and horizontal flipping for training, resizing and center cropping for  
 731 evaluation).

## 732 H Training Metrics

733 In the Fig. 7 below we monitor the evolution of both the objective (second panel), the mean  
 734 augmentation manifold nuclear norm, the centroid norm, and the mean centroid similarity evaluated  
 735 on the test set over the course of training.




Figure 7: Evolution of various metrics during training. Geometric measures are evaluated on a set of 200 manifolds, each defined by an image drawn from the CIFAR-10 dataset, along with 16 augmentations. Shaded regions indicate a 95% confidence interval around the mean.

## 736 I Classification Performance on Smaller Datasets

737 In Table 4 below we report the performance of both our method as well as Barlow Twins and SimCLR  
 738 when trained using a ResNet-50 backbone on smaller datasets.

## 739 J Batch Size Dependence

740 One of the most cited drawbacks of contrastive SSL methods has been that strong performance  
 741 on downstream tasks requires training with large batch sizes, while non-contrastive methods (e.g.,  
 742 VICReg or Barlow Twins (63; 6)) that place constraints on the cross-correlation/covariance matrices  
 743 of the embeddings are much more amenable to smaller batch training. It is also worth noting that  
 744 the need for large batch sizes in contrastive methods can be alleviated in various ways, such as  
 745 maintaining a memory bank (61) or employing a slowly updating momentum encoder (36). Given  
 746 that our method is neither wholly contrastive nor non-contrastive (since it acts on the spectrum of the  
 747 embedding matrix directly), we wondered how would depend on training batch size. We pretrained  
 748 on ImageNet-1k using batch sizes of 256, 512, 1024, 2048, 4096 and evaluate the linear classification

Table 4: Top-1 classification accuracies of linear classifiers for representations trained with various datasets and objective functions. Note: for Barlow Twins on ImageNet-100 we report the result from da Costa et al. (20) which uses a ResNet-18 backbone, as we were unable to obtain better performance. For MMCR on ImageNet-100 we tested both 2 views (matched to baselines) and 4 views, results are formatted (2-view)/(4-view)

	Method	CIFAR-10	CIFAR-100	STL-10	ImageNet-100
[t]	Barlow Twins (our repro.)	90.91	67.91	89.96	80.38*
	SimCLR (our repro.)	92.22	70.04	91.11	79.64
	MMCR ( $\lambda = 0.0$ )	93.53	69.87	90.62	81.52/82.88
	MMCR ( $\lambda = 0.01$ )	93.39	70.94	90.77	81.28/82.56

749 accuracy for each. Encouragingly we observed only a modest decrease in performance for the  
 750 smallest batch size tested. The results of this sweep, in comparison to Barlow Twins and SimCLR, is  
 751 shown in in Fig. 8 Note that for these runs we used two views and the linear learning rate scaling as  
 752 described in Appendix D. Future work should endeavor to better understand the impact of various  
 753 hyperparameters on the quality of learned representations.

754 An important detail is that for this experiment we did not employ a momentum encoder when training  
 755 MMCR. It is argued in He et al. (36) that the momentum encoder increases the effective minibatch  
 756 size as the slowly moving weights encode information from preceding batches. However here we  
 757 are interested explicitly interested in batch size dependence so we ablate this architectural confound  
 758 (neither Zbontar et al. (63); Chen et al. (12) employ momentum encoders).




Figure 8: Drop in top-1 performance relative to that of the best setting for three methods. Data for both Barlow Twins and SimCLR are copied from Zbontar et al. (63).

## 759 K Additional Details on BrainScore

760 Brain-Score evaluates a model in terms of its ability to predict the measured responses of neurons to  
 761 images. We provide a brief introduction to the metric here (see (54) for a complete description). Let  
 762  $\mathbf{y} \in \mathbb{R}^N$  denote the average response of a single (biological) neuron to a set of  $N$  training images  
 763 and  $\mathbf{X} \in \mathbb{R}^{N \times K}$  be the response of  $K$  model neurons to the same images. Brain Score first solves  
 764 the linear regression problem  $\mathbf{y} = \mathbf{X}\mathbf{w}$  for weights  $\mathbf{w}$ . The model then predicts responses  $\mathbf{y}'$  to a set  
 765 of held out images. Next the Pearson correlation coefficient between  $\mathbf{y}'$  and  $\mathbf{y}$  is calculated, and the  
 766 score for a particular dataset is the median of the individual neuron predictivities in said dataset. The  
 767 mean of these scores is taken over different train-test splits of each dataset. The score for a total brain  
 768 area (as shown in Table 2 is the mean over train-test splits and distinct datasets. The standard error of  
 769 the means for each dataset (which are typically between 1e-3 and 1e-2) within an area are summed  
 770 quadrature and divided by the number of datasets to produce the errors reported in table 2. In table  
 771 5 we split the brain area scores into individual datasets.

Model	V1.0	V1.1	V2.0	V4.0	V4.1	V4.2
MMCR	0.270	0.718	<u>0.311</u>	<u>0.577</u>	<u>0.627</u>	<u>0.492</u>
SimCLR	0.224	<u>0.776</u>	0.288	0.576	0.626	0.48
BYOL	<u>0.274</u>	<u>0.727</u>	0.291	<u>0.585</u>	0.626	0.48
MoCo	0.273	.726	0.293	0.57	<u>0.629</u>	<u>0.492</u>
Barlow	<u>0.276</u>	.721	0.293	0.568	0.626	<u>0.493</u>
SwAV	0.252	.723	<u>0.296</u>	0.568	0.614	0.469
Model	V4.3	IT.0	IT.1	IT.2	IT.3	
MMCR	0.226	<u>0.554</u>	<u>0.558</u>	<u>0.545</u>	0.424	
SimCLR	<u>0.224</u>	<u>0.552</u>	<u>0.545</u>	<u>0.518</u>	<u>0.456</u>	
BYOL	0.216	0.55	0.545	0.516	0.41	
MoCo	0.215	0.54	<u>0.560</u>	<u>0.550</u>	<u>0.437</u>	
Barlow	0.221	0.545	0.547	0.518	0.412	
SwAV	0.202	0.533	0.537	0.518	0.405	

Table 5: Brain-score comparison of six self-supervised models, over 11 different electrophysiological data sets recorded from macaque monkeys (54). Datasets V1.0 and V2.0 are from Freeman et al. (26), V1.1 is from Marques et al. (47), and V4.0, V4.1, IT.0, and IT.1 are from Majaj et al. (45). (V4/IT).(2/3) are the SanghaviJozwik2020 and SanghaviMurty2020 datasets as denoted by brainscore.

## 772 L Additional Details on Spectral Properties

773 **Participation Ratio** We first extracted the 2048 dimensional feature vectors for each model in  
774 response to the images in the ImageNet validation set. Images were resized to  $256 \times 256$  and then  
775 center cropped to  $224 \times 224$  following the setting in which the classifiers are tested. We resampled  
776 the resultant feature matrices with replacement 10 times independently for each model. For each  
777 resampled dataset (of features) we calculate the empirical covariance matrix, associated eigenspectra,  
778 and associated participation ratios (squared ratio of  $L_1$  to  $L_2$  norm of the eigenvectors).

779 **Decay Coefficient** The spectra obtained using the procedure outlined above all decayed rapidly near  
780 the tails (the least significant eigenvalues). To avoid undo bias from these tails when estimating the  
781 decay coefficients we only considered the top 2000 eigenvalues. To estimate the decay coefficient we  
782 fit a regression line to the logarithm of the eigenvalues as a function of the logarithm of their indexes.  
783 This fitting procedure was repeated across the bootstrapped spectra for each model to obtain standard  
784 errors of the mean. For all models the linear regression produced a strong fit to the data, with the  
785 minimum observed  $R^2$  value being 0.95.

## 786 M Object Detection

787 To ensure that the representations obtained with MMCR are not hyperspecialized to classification  
788 tasks we also evaluated our highest performing model on object detection. We follow (36; 63), fine  
789 tuning the representation network with a Faster R-CNN head and C-4 backbone on the VOC07+12  
790 dataset (training with the train+val split and evaluating on the VOC07 test split). All settings except  
791 for the initial learning rate (which we set to 0.12) were identical to those from He et al. (36), and  
792 we similarly used the detectron2 library for this evaluation. A representation trained using MoCo  
793 v2 for 200 epochs achieves an AP50 (the most common evaluation metric for this dataset) of 82.4.  
794 MMCR with 8 views and 100 epochs of pretrained produced an AP50 of 81.9 (this is the mean over  
795 three independent fine tunes, the standard deviation was 0.2), demonstrating that the representations  
796 generated by our method are not limited to object recognition.

## 797 N Limitations

798 Time and compute limitations prevented us from conducting exhaustive optimization of all design  
799 choices. For example we do not explore effects of varying the projector network width and depth.  
800 We additionally restrict the model to a ResNet-50 encoding network pretrained on the ImageNet-1k  
801 dataset for 100 epochs. This choice allows us to make fair comparisons to several recently developed

802 SSL methods with a modest compute budget, but precludes answering questions about how the  
803 method scales up to larger problems.

804 We also note that although our objective function has favorable computational complexity compared to  
805 existing methods, the evaluation of our objective is not straightforward to distribute across machines.  
806 This is because we need to compute the SVD over the centroid matrix, which requires gathering the  
807 outputs of a distributed forward pass onto a single machine. Efficient methods for the distributed  
808 computation of the SVD could help alleviate this issue in the future.