
Understanding self-supervised Learning Dynamics without Contrastive Pairs

Yuandong Tian¹ Xinlei Chen¹ Surya Ganguli^{1 2}

Abstract

Contrastive approaches to self-supervised learning (SSL) learn representations by minimizing the distance between two augmented views of the same data point (positive pairs) and maximizing the same from different data points (negative pairs). However, recent approaches like BYOL and SimSiam, show remarkable performance *without* negative pairs, raising a fundamental theoretical question: how can SSL with only positive pairs avoid representational collapse? We study the nonlinear learning dynamics of non-contrastive SSL in simple linear networks. Our analysis yields conceptual insights into how non-contrastive SSL methods learn, how they avoid representational collapse, and how multiple factors, like predictor networks, stop-gradients, exponential moving averages, and weight decay all come into play. Our simple theory recapitulates the results of real-world ablation studies in both STL-10 and ImageNet. Furthermore, motivated by our theory we propose a novel approach that *directly* sets the predictor based on the statistics of its inputs. In the case of linear predictors, our approach outperforms gradient training of the predictor by 5% and on ImageNet it performs comparably with more complex two-layer non-linear predictors that employ BatchNorm. Code is released in <https://github.com/facebookresearch/luckmatters/tree/master/ssl>.

1. Introduction

Self-supervised learning (SSL) has emerged as a powerful method for learning useful representations without requiring expensive target labels. Many SSL methods employ the principle of contrastive learning (Oord et al., 2018; Tian

¹Facebook AI Research ²Stanford University. Correspondence to: Yuandong Tian <yuandong@fb.com>.

et al., 2019; He et al., 2020; Chen et al., 2020a; Bachman et al., 2019) whereby the hidden representations of two augmented views of the same object (positive pairs) are brought closer together, while those of different objects (negative pairs) are encouraged to be further apart. Minimizing differences between positive pairs encourages modeling invariances, while contrasting negative pairs is thought to be required to prevent representational collapse (i.e., mapping all data to the same representation)

However, recent SSL work, notably BYOL (Grill et al., 2020) and SimSiam (Chen & He, 2020) have shown the remarkable capacity to learn powerful representations using only positive pairs, *without* ever contrasting negative pairs. These methods employ a dual pair of Siamese networks (Bromley et al., 1994) (Fig. 1): the representation of two views are trained to match, one obtained by the composition of an online and predictor network, and the other by a target network. The target network is *not* trained via gradient descent; it instead acts as a momentum encoder (see e.g., MoCo (He et al., 2020; Chen et al., 2020b), PIRL (Misra & Maaten, 2020), SwAV (Caron et al., 2020), and BYOL (Grill et al., 2020)) that slowly follows the online network in a delayed fashion through an exponential moving average (EMA) computation.

Since this entire procedure encourages the online+predictor network and the target network to become similar to each other, this overall scheme raises several fundamental unsolved theoretical questions. Why/how does it avoid collapsed representations? What is the nature of the learned representations? How do multiple design choices and hyperparameters interact nonlinearly in the learning dynamics? While there are interesting theoretical studies of contrastive SSL (Arora et al., 2019; Lee et al., 2020; Tosh et al., 2020), any theoretical understanding of the nonlinear learning dynamics of non-contrastive SSL remains open.

In this paper, we make a first attempt to analyze the behavior of non-contrastive SSL training and the empirical effects of multiple hyperparameters, including (1) Exponential Moving Average (EMA) or momentum encoder, (2) Higher relative learning rate (α_p) of the predictor, and (3) Weight decay η . We explain all these empirical findings with an exceedingly simple theory based on analyzing the nonlinear learning dynamics of simple linear networks.

Note that deep linear networks have provided a useful tractable theoretical model of nonconvex loss landscapes (Kawaguchi, 2016; Du & Hu, 2019; Laurent & Brecht, 2018) and nonlinear learning dynamics (Saxe et al., 2013; 2019; Lampinen & Ganguli, 2018; Arora et al., 2018) in these landscapes, yielding insights like dynamical isometry (Saxe et al., 2013; Pennington et al., 2017; 2018) that lead to improved training of nonlinear deep networks. Despite the simplicity of our theory, it can still predict how various hyperparameters choices affect performance in an extensive set of real-world ablation studies. Moreover, the simplicity also enables us to provide conceptual and analytic insights into *why* performance patterns vary the way they do. Specifically, our theory accounts for the following diverse empirical findings:

Essential part of non-contrastive SSL. The existence of the predictor and stop-gradient is absolutely essential. Removing either of them leads to representational collapse in BYOL and SimSiam.

EMA. While the original BYOL needs EMA to work, they later confirmed that EMA is not necessary (i.e., the online and target networks can be identical) if a higher α_p is used. This is also confirmed with SimSiam, as long as the predictor is updated more often or has larger learning rate (or larger α_p). However, the performance is slightly lower.

Predictor Optimality and Relative learning rate α_p . Both BYOL and SimSiam suggest that the predictor should always be optimal, in the sense of always achieving minimal ℓ_2 error in predicting the target network’s outputs from the online network’s outputs. This optimality conjecture was motivated by observed superior performance when the predictor had large learning rates and/or was allowed more frequent updates than the rest of the network. However (Chen & He, 2020) also showed that if the predictor is updated too often, then performance drops, which questions the importance of an always optimal predictor as a key requirement for learning good representations.

Weight Decay. Table 15 in BYOL (Grill et al., 2020) indicates that no weight decay may lead to unstable results. A recent blogpost (Fetterman & Albrecht, 2020) also mentions using weight decay leads to stable learning in BYOL.

Finally, motivated by our theoretical analysis, we propose a new method **DirectPred** that directly sets the predictor weights based on principal components analysis of the predictor’s input, thereby avoiding complicated predictor dynamics and initialization issues. We show that this simple **DirectPred** method nevertheless yields comparable performance in CIFAR-10 and outperforms gradient training of the linear predictor by +5% Top-1 accuracy in linear evaluation protocol on both STL-10 and ImageNet (60 epochs).

	Plug-in frequency (every N minibatches)			
	1	2	3	5
EMA	40.67 ± 0.50	35.29 ± 2.49	34.60 ± 0.98	35.63 ± 2.66
no EMA	39.45 ± 1.26	34.01 ± 1.54	34.58 ± 2.93	32.22 ± 2.94

Table 1. Simply plugging in the “optimal solution” to the linear predictor shows poor performance after 100 BYOL epochs (Top-1 accuracy in STL-10 (Coates et al., 2011) downstream classification task). The optimal solution is obtained by solving (with regularization) $W_p \mathbb{E}[\mathbf{f} \mathbf{f}^\top] = \frac{1}{2} (\mathbb{E}[\mathbf{f}_a \mathbf{f}^\top] + \mathbb{E}[\mathbf{f} \mathbf{f}_a^\top])$, in which the two expectations is estimated with exponential moving average. In comparison, with gradient descent, BYOL with a single linear layer predictor can reach 74%-75% Top-1 in STL-10 after 100 epochs. Unless explicitly stated, in all our experiments, we use ResNet-18 (He et al., 2016) as the backbone network and SGD as the optimizer with learning rate $\alpha = 0.03$, momentum 0.9, weight decay $\bar{\eta} = 0.0004$ and EMA parameter $\gamma_a = 0.996$. Each setting is repeated 5 times.




Figure 1. Two-layer setting with a linear, bias-free predictor.

2. Two-layer linear model

To obtain analytic and conceptual insights into non-contrastive SSL we analyze a simple, *bias-free* linear BYOL model where the online, target and predictor networks are specified by the weight matrices W , W_p and W_a respectively (Fig. 1). Let $\mathbf{x} \in \mathbb{R}^{n_1}$ be a data point drawn from the data distribution $P(\mathbf{x})$ and let \mathbf{x}_1 and \mathbf{x}_2 be two augmented views of \mathbf{x} : $\mathbf{x}_1, \mathbf{x}_2 \sim p_{\text{aug}}(\cdot|\mathbf{x})$ where $p_{\text{aug}}(\cdot|\mathbf{x})$ is the augmentation distribution. In practice such data augmentations correspond to random crops, blurs or color distortions of images (Chen et al., 2020a). Let $\mathbf{f}_1 = W\mathbf{x}_1 \in \mathbb{R}^{n_2}$ be the online representation of view 1, and $\mathbf{f}_{2a} = W_a\mathbf{x}_2 \in \mathbb{R}^{n_2}$ be the target representation of view 2. In BYOL, the learning dynamics of W and W_p are obtained by minimizing

$$J(W, W_p) := \frac{1}{2} \mathbb{E}_{\mathbf{x}_1, \mathbf{x}_2} [\|W_p \mathbf{f}_1 - \text{StopGrad}(\mathbf{f}_{2a})\|_2^2], \quad (1)$$

while the dynamics of W_a is obtained differently, via an exponential moving average (EMA) of W . We will analyze this combined dynamics for W , W_p and W_a , in the pres-

ence of additional weight decay, in the limit of large batch sizes and small discrete time learning rates. This limit can be well approximated by the gradient flow (see Supplementary Material (SM) for all derivations):

Lemma 1. *BYOL learning dynamics following Eqn. 1:*

$$\dot{W}_p = \alpha_p (-W_p W (X + X') + W_a X) W^\top - \eta W_p \quad (2)$$

$$\dot{W} = W_p^\top (-W_p W (X + X') + W_a X) - \eta W \quad (3)$$

$$\dot{W}_a = \beta (-W_a + W) \quad (4)$$

Here, $X = \mathbb{E}[\bar{x}\bar{x}^\top]$ where $\bar{x}(x) := \mathbb{E}_{x' \sim p_{\text{aug}}(\cdot|x)}[x']$ is the average augmented view of a data point x and $X' = \mathbb{E}_x[\mathbb{V}_{x'|x}[x']]$ is the covariance matrix $\mathbb{V}_{x'|x}[x']$ of augmented views x' conditioned on x , subsequently averaged over the data x . Note that α_p and β reflect *multiplicative learning rate ratios* between the predictor and target networks relative to the online network. Finally, the terms involving η reflect weight decay.

As a gradient flow formulation, the learning rate α does not appear in Lemma 1. In the actual finite time update, the learning rate for W_p is $\alpha\alpha_p$, the EMA rate is $\alpha\beta = 1 - \gamma_a$, where γ_a is the usual EMA parameter (e.g., BYOL uses 0.996), and the weight decay for actual training is $\tilde{\eta} := \alpha\eta$.

We note that since SimSiam is an ablation of BYOL that removes the EMA computation, the underlying dynamics of SimSiam can also be obtained from Lemma 1 simply by setting $W_a = W$, inserting this relation into Eqn. 2 and Eqn. 3, and ignoring Eqn. 4. Importantly, the stop-gradient on the target branch is still there.

Overall Eqns. 2-4 constitute our starting point for analyzing the combined roles of relative learning rates α_p and β , weight decay rate η and various ablations in determining the performance of both BYOL and SimSiam.

We first derive two very general results (see SM).

Theorem 1 (Weight decay promotes balancing of the predictor and online networks.). *Completely independent of the particular dynamics of W_a in Eqn. 4, the update rules (Eqn. 2 and Eqn. 3) possess the invariance*

$$W(t)W^\top(t) = \alpha_p^{-1}W_p^\top(t)W_p(t) + e^{-2\eta t}C, \quad (5)$$

where C is a symmetric matrix that depends only on the initialization of W and W_p .

This theorem implies that for both BYOL and SimSiam, there exists a ‘‘balancing’’ that ensures that any matching between the online and target representations will not be attributable solely to the predictor weights, rendering the online weights useless. Instead what the predictor learns, the online network will also learn, which is important as the online network’s representations are what is used for

EMA + no-bias	EMA + bias	no EMA + no-bias	no EMA + bias
70.62 ± 1.05	70.99 ± 1.01	71.36 ± 0.44	71.37 ± 0.77

Table 2. Top-1 accuracy of BYOL on STL-10 under linear evaluation protocol, trained for 100 epochs with no weight decay ($\eta = 0$) and $\alpha_p = 1$. It is worse than the baseline (74.51 ± 0.47 without predictor bias) when the weight decay is set to be $\eta = 0.0004$. ‘‘No-bias’’ means the linear predictor does not have a bias term.

downstream tasks. We note that similar weight balancing dynamics has been discovered in multi-layer linear networks and matrix factorization (Arora et al., 2018; Du et al., 2018). Our results generalize this to SSL dynamics. Second, a nonzero weight decay could help remove the extra constant C due to initialization, further balancing the predictor and online network weights and possibly leading to better performance on downstream tasks (Tbl. 2).

Theorem 2 (The stop-gradient signal is essential for success.). *With $W_a = W$, removing the stop-gradient signal yields a gradient update for W given by*

$$\frac{d}{dt}\text{vec}(W) = - \left[X' \otimes (W_p^\top W_p + I) + X \otimes \tilde{W}_p^\top \tilde{W}_p \right] \text{vec}(W). \quad (6)$$

Here $\tilde{W}_p = W_p - I$ and \otimes is the Kronecker product. This dynamics leads to $W(t) \rightarrow 0$ when X' is full-rank.

Thus we have proven analytically in this simple setting that removing the stop-gradient leads to representational collapse, as observed in more complex settings in SimSiam (Chen & He, 2020).

3. How multiple factors affect learning dynamics

The learning dynamics in Eqns. 2-4 constitute a set of high dimensional coupled nonlinear differential equations that can be difficult to solve analytically in general. Therefore, to obtain analytic insights into the functional roles of the relative learning rates α_p and β and weight decay η , we make a series of simplifying assumptions. Intriguingly, under these simplifying assumptions we obtain a rich set of analytic predictions, which we then test experimentally in more realistic scenarios. We find, nicely, that these predictions still qualitatively hold *even when* our simplifying assumptions required for obtaining analytic results do not.

Assumption 1 (Proportional EMA). *We first reduce the dimensionality of the dynamics in Eqns. 2-4 by enforcing that the target network W_a undergoes EMA but is forced to always be proportional to the online network via the relation $W_a(t) = \tau(t)W(t)$. Inserting this relation into the EMA dynamics in Eqn. 4 yields $\dot{W} + \tau\dot{W} = \beta(1 - \tau)W$.*

Thus we obtain a reduced dynamics for W , W_p and τ . By not enforcing the stronger SimSiam constraint that

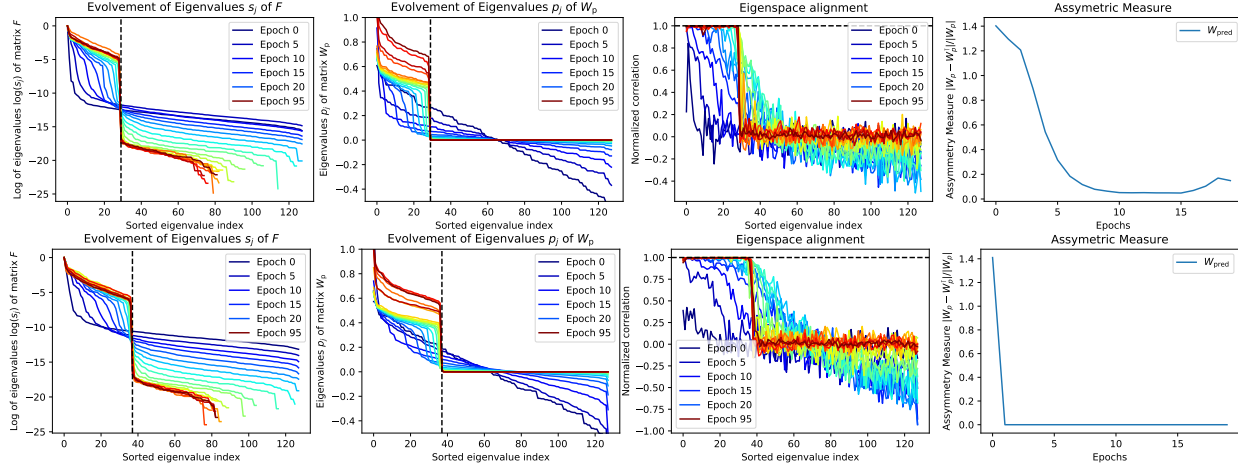


Figure 2. Training BYOL in STL-10 for 100 epochs with EMA. **Top row**: No symmetric regularization imposed on W_p , **Bottom row**: symmetric regularization on W_p . From left to right: **(1)** Evolution of eigenvalues for F . Since F is PSD and its eigenvalue s_j varies across scales, we plot $\log(s_j)$. We could see some eigenvalues are growing while others are shrinking to zero over training. **(2)** Similar “step-function” behaviors for the predictor W_p . Its negative eigenvalues shrinks towards zero and leading eigenvalues becomes larger. **(3)** The eigenspace of F and W_p gradually align with each other (Theorem 3). For each eigenvector u_j of F , we compute cosine angle (normalized correlation) between u_j and $W_p u_j$ to measure alignment. **(4)** W_p gradually becomes symmetric and PSD during training.

$W_a = W$, we can still model EMA dynamics. Intuitively, $\tau = \tau(t)$ is a dynamic parameter that depends on how quickly $W = W(t)$ grows over time. If W is constant, then $\dot{W} = 0$ and τ stabilizes to 1. On the otherhand, if W grows rapidly, then τ becomes small. While Assumption 1 is a simplification, as we shall see, it still reveals interesting verifiable predictions about the functional role of EMA.

Assumption 2 (Isotropic data and augmentation). *We assume the data distribution $p(x)$ has zero mean and identity covariance, while the augmentation distribution $p_{\text{aug}}(\cdot|x)$ has mean x and covariance $\sigma^2 I$. This simplifies the dynamics in Eqns. 2-4 by reducing the augmentation averaged data covariance to $X = I$ and the data averaged augmentation covariance to $X' = \sigma^2 I$.*

Many previous studies of deep learning dynamics made simplifying isotropic assumptions about data (Tian, 2017; Brutzkus & Globerson, 2017; Du et al., 2019; Bartlett et al., 2018; Safran & Shamir, 2018). Since our fundamental goal is to obtain the first analytic understanding of the dynamics of non-contrastive SSL methods, it is useful to first achieve this in the simplest possible isotropic setting. Interestingly, we will find that our final conclusions generalize to non-isotropic real world settings.

Assumption 3 (Symmetric predictor). *We enforce symmetry in W_p by initializing it to be a symmetric matrix, and then symmetrizing the flow for W_p in Eqn. 2 (see SM).*

This symmetry assumption was motivated by both fixed point analysis and empirical findings. First, the fixed point of Eqn. 2 under Assumption 1 and 2 and $\eta > 0$ is always a symmetric matrix and in numerical simulation the asym-

metric part $W_p - W_p^\top$ eventually vanishes (See Appendix for the proof and numerical simulations). Moreover, during BYOL training without a symmetry constraint on the predictor, W_p gradually moves towards symmetry (Fig. 2).

Second, a set of experiments reveal that whether the predictor is symmetric or not has a dramatic effect in terms of both performance and interaction with EMA. In our STL-10 experiment, enforcing symmetric W_p in the presence of EMA *improves* performance on downstream tasks (Tbl. 3). In contrast, in the absence of EMA, a symmetric W_p fails while an asymmetric W_p works reasonably well. Similar behavior holds on ImageNet: a symmetric one layer linear predictor W_p in SimSiam (i.e. without EMA) achieves performance no better than random guessing (Top-1/5: 0.1%/0.5%), while an asymmetric W_p achieves a Top-1/5 accuracy of 68.1%/88.2%. Our theory will explain this as well as show how to obtain good performance with a symmetric predictor without EMA by increasing its relative learning rate α_p .

3.1. Dynamical alignment of eigenspaces between the predictor and its input correlation matrix

Under the 3 assumptions stated above, we analyze the coupled dynamics of $F = W X W^\top$ and W_p . Note that F is the correlation matrix of the outputs of the online network which also serve as inputs to the predictor. We find F and W_p obey the following dynamics (see SM):

$$\begin{aligned} \dot{W}_p &= -\frac{\alpha_p}{2}(1 + \sigma^2)\{W_p, F\} + \alpha_p \tau F - \eta W_p \quad (7) \\ \dot{F} &= -(1 + \sigma^2)\{W_p^2, F\} + \tau\{W_p, F\} - 2\eta F \end{aligned}$$

	No predictor bias		With predictor bias	
	sym W_p	regular W_p	sym W_p	regular W_p
<i>One-layer linear predictor</i>				
EMA	75.09 ± 0.48	74.51 ± 0.47	74.52 ± 0.29	74.16 ± 0.33
no EMA	36.62 ± 1.85	72.85 ± 0.16	36.04 ± 2.74	72.13 ± 0.53
<i>Two-layer predictor with BatchNorm and ReLU</i>				
EMA	71.58 ± 6.46	78.85 ± 0.25	77.64 ± 0.41	78.53 ± 0.34
no EMA	35.59 ± 2.10	65.98 ± 0.71	41.92 ± 4.25	65.59 ± 0.66

Table 3. The effect of symmetrization of W_p on downstream classification task (BYOL Top-1 on STL-10). Symmetric W_p leads to slightly better performance compared to regular W_p in the presence of EMA. On the other hand, without EMA, symmetric W_p crashes. Same effects happen in two-layer predictor with BatchNorm and ReLU as well. Weight decay $\bar{\eta} = 0.0004$ and $\alpha_p = 1$.

This dynamics reveals that the eigenspace of W_p will gradually align with that of F under certain conditions (see SM for derivation):

Theorem 3 (Eigenspace alignment). *Under Eqn. 7, the commutator $[F, W_p] := FW_p - W_pF$ satisfies:*

$$\frac{d}{dt}[F, W_p] = -[F, W_p]K - K[F, W_p] \quad (8)$$

where

$$K(t) = (1 + \sigma^2) \left[\frac{\alpha_p}{2} F(t) + W_p^2(t) - \frac{\tau}{1 + \sigma^2} W_p(t) \right] + \frac{3}{2} \eta I \quad (9)$$

If $\max_{t \geq 0} \lambda_{\min}[K(t)] = \lambda_0 > 0$, then the commutator

$$\|[F(t), W_p(t)]\|_F \leq e^{-2\lambda_0 t} \|[F(0), W_p(0)]\|_F \rightarrow 0 \quad (10)$$

For symmetric W_p , when W_p and F commute they can be simultaneously diagonalized. Thus this shows that the eigenspace of W_p gradually aligns with that of F .

To test this prediction, we performed extensive experiments showing that training BYOL using ResNet-18 on STL-10 yields eigenspace alignment, as demonstrated in Fig. 2.

Now if the eigenspaces of W_p and F do align, we can obtain fully decoupled dynamics. Let the columns of the matrix U be the common eigenvectors, so that $W_p = U\Lambda_{W_p}U^\top$ where $\Lambda_{W_p} = \text{diag}[p_1, p_2, \dots, p_d]$, $F = U\Lambda_F U^\top$ where $\Lambda_F = \text{diag}[s_1, s_2, \dots, s_d]$. For each mode j , we have (see SM for derivation):

$$\dot{p}_j = \alpha_p s_j [\tau - (1 + \sigma^2)p_j] - \eta p_j \quad (11)$$

$$\dot{s}_j = 2p_j s_j [\tau - (1 + \sigma^2)p_j] - 2\eta s_j \quad (12)$$

$$s_j \dot{\tau} = \beta(1 - \tau)s_j - \tau \dot{s}_j / 2. \quad (13)$$

This decoupled dynamics constitutes a dramatically simplified set of 3 dimensional nonlinear dynamical systems for BYOL learning, and two dimensional nonlinear systems (obtained by constraining $\tau = 1$) for SimSiam. As

expected, each mode’s dynamics is equivalent to the 3 dimensional dynamics obtained by setting $n_1 = n_2 = 1$ in Eqns. 2-4 and making the replacements $W^2 = s_j$, $W_p = p_j$, and $W_a/W = \tau$ (see SM). Thus the decoupled dynamics in Eqns 11- 13 reduce to the scalar case of BYOL dynamics in Eqns. 2-4 after a change of variables and the condition in Thm. 3 reveals when this decoupled regime is reachable.

3.2. Analysis of decoupled dynamics

The simplified three (two) dimensional dynamics of BYOL (SimSiam) yields significant insights. First, there is clearly a collapsed fixed point at $p_j(t) = s_j(t) = 0$ and τ taking any value. We wish to understand conditions under which p_j and s_j can avoid this collapsed fixed point and grow from small random initial conditions. Since s_j is an eigenvalue of WW^\top , we are particularly interested in conditions under which s_j achieves large final values, corresponding to a non-collapsed online network, that are moreover sensitive to the statistics of the data, governed by σ^2 .

Exact integral. First, an important observation, similar to Theorem 1, is that the dynamics possesses an exact integral of motion, obtained by multiplying Eqn. 11 by $2\alpha_p^{-1}p_j$, subtracting, Eqn. 12 and integrating over time yielding

$$s_j(t) = \alpha_p^{-1}p_j^2(t) + e^{-2\eta t}c_j \quad (14)$$

where $c_j = \alpha_p^{-1}p_j^2(0) - s_j(0)$ is fixed by initial conditions. In absence of weight decay ($\eta = 0$), this integral reveals that the initial condition encoded in c_j is never forgotten and the dynamics of p_j and s_j are confined to parabolas of the form $s_j(t) = p_j^2(t) + c_j$, as can be seen by the blue flow lines in Fig. 3(left). With weight decay ($\eta > 0$) over time the initial condition is forgotten and the dynamics approaches the invariant parabola $s_j = \alpha_p^{-1}p_j^2$ as can be seen by the approach of the blue flow lines to the black dashed parabola in Fig. 3 right and middle. We discuss these two cases in turn. First we note that in both cases, since the EMA computation is often very slow (Grill et al., 2020), corresponding to small β , the dynamics of τ in Eqn. 13 is slow relative to that of p_j and s_j . Therefore to understand the combined dynamics, we can search for the fixed points that p_j and s_j will rapidly approach at fixed τ . Over time τ will then either slowly approach 1 (BYOL) or be always equal to 1 (SimSiam), and s_j and p_j will follow their τ -dependent fixed points.

No weight decay. When $\eta = 0$, Eqns. 11 and 12 at a fixed value of τ yield a branch of collapsed fixed points given by $s_j = 0$ and p_j taking any value, and a branch of non-collapsed fixed points, with $p_j = \tau/(1 + \sigma^2)$ and s_j taking any value (horizontal and vertical red/green lines in Fig. 3,left). A sufficient criterion on initial conditions to avoid the collapsed branch is $s_j(0) > p_j^2(0)/\alpha_p$ cor-




Figure 3. State space dynamics in Eqns. 11 and 12 for no ($\eta = 0$) weak ($\eta = 0.01$) and strong ($\eta = 1$) weight decay at fixed $\tau = 1$ and $\alpha_p = 1$. Red (green) points indicate stable (unstable) fixed points, blue curves indicate flow lines, and the dashed black curve indicates the parabola $s_j = p_j^2 / \alpha_p$.

	Positive effects	Negative effects
Relative predictor lr α_p	#1,#6	#2
Weight decay η	#3,#7	#4,#5
EMA β	#8	#9,#10

Table 4. Summarization of positive/negative effects of various hyperparameter choices (EMA β , relative predictor learning rate α_p and weight decay η). “#1” means (Obs#1) in the text.

responding to lying above the dashed black parabola in Fig. 3, left. This restricted initial condition reveals why a fast predictor (large α_p) is advantageous (Obs#1): larger α_p leads to a smaller basin of attraction of the collapsed branch by flattening the dashed parabola. Indeed both BYOL and SimSiam have noted that a fast predictor can help avoid collapse. On the other hand, α_p cannot be infinitely large (Obs#2): since $s_j(+\infty) = s_j(0) + \alpha_p^{-1}(p_j^2(+\infty) - p_j^2(0))$, very large α_p implies that s_j , the final value of the online network characterizing the learned representation, does not grow even if p_j does. This is consistent with results which show that optimizing the predictor too often doesn’t work in SimSiam (Chen & He, 2020), and directly setting an “optimal” predictor fails as well (Tbl. 1). The online network needs to grow along with the predictor and that cannot happen if the predictor is too fast.

Advantage of weight decay. In the non-collapsed branch of fixed points without weight decay (vertical red line in Fig. 3, left), the predictor p_j takes the exact value $\tau / (1 + \sigma^2)$, which models the invariance to augmentation correctly: a large data augmentation variance σ^2 should lead to a small magnitude of the learned representation. Ideally, we want s_j to have the same property. With weight decay $\eta > 0$ in Eqn. 14, memory of the initial condition c_j fades away, yielding convergence to some point on the invariant parabola $s_j = \alpha_p^{-1} p_j^2$. (Obs#3): Therefore, by tying the online network to the predictor, weight decay allows s_j to also model invariance to augmentations correctly if the predictor does, regardless of the random initial condition c_j .

Dynamics on the invariant parabola. Because weight decay forces convergence to the invariant parabola $s_j = \alpha_p^{-1} p_j^2$, we next focus on dynamics along this parabola (i.e.




Figure 4. Fixed point of $\dot{p}_j = p_j(p_j - p_{j-}^*)(p_j - p_{j+}^*)$. Stable fixed points are in red, unstable in green and saddle in black. When the weight decay $\eta = 0$, the trivial solution $p_j = 0$ is a saddle. When $\eta > 0$, the trivial solution becomes stable near to the origin and initial p_j needs to be large enough to converge to the stable non-collapsed solution p_{j+}^* .

$c_j = 0$ in Eqn. 14). In this case, Eqn. 13 has a solution:

$$\tau(t) = p_j^{-1}(t) \beta e^{-\beta t} \int_0^t p_j(t') e^{\beta t'} dt, \quad (15)$$

with initial condition $\tau(0) = 0$. Inserting the invariant $s_j = \alpha_p^{-1} p_j^2$ into Eqn. 11, the dynamics of p_j is given by:

$$\dot{p}_j = p_j^2 [\tau(t) - (1 + \sigma^2)p_j] - \eta p_j. \quad (16)$$

We first analyze the fixed points where $\dot{p}_j = 0$ at fixed τ . When the weight decay $0 < \eta \leq \frac{\tau^2}{4(1 + \sigma^2)}$, p_j has three fixed points (Fig. 4(b)):

$$p_{j\pm}^* = \frac{\tau \pm \sqrt{\tau^2 - 4\eta(1 + \sigma^2)}}{2(1 + \sigma^2)} > 0, \quad p_{j0}^* = 0$$

where both p_{j0}^* and p_{j+}^* are stable and p_{j-}^* is unstable, as shown in Fig. 4(b). The basin of attraction of the collapsed fixed point $p_{j0}^* = 0$ is $p_j < p_{j-}^*$ while the basin of attraction of the useful non-collapsed fixed point p_{j+}^* is $p_j > p_{j-}^*$, yielding an important constraint on initial conditions to avoid collapse. Note that p_{j-}^* is a decreasing function of τ and increasing function of η (see SM). This means that with larger η , p_{j-}^* moves right and the basin of collapse expands (Obs#4). When $\eta > \frac{\tau^2}{4(1 + \sigma^2)}$ there is only one stable fixed point $p_{j0}^* = 0$ (Fig. 4(c)). Under such strong weight decay collapse is unavoidable (Obs#5).

We now discuss the dynamics. First we define the quantity $\Delta_j := p_j[\tau - (1 + \sigma^2)p_j] - \eta$, which must satisfy *two criteria*. Note that Eqn. 16 can be written as $\dot{p}_j = p_j \Delta_j$, so Δ_j must at some point be positive to drive $p_j(t)$ to any positive non-collapsed fixed point p_{j+}^* . Second, for eigenspace alignment in Theorem 3 to *remain* stable (even if the alignment has already happened), $K(t)$ must be positive definite (PD) in Eqn. 9. Using the eigenspace alignment conditions and the invariance $s_j = \alpha_p^{-1} p_j^2$, the positive definite condition on $K(t)$ can be written as

$$\Delta_j < \frac{1}{2} [\alpha_p(1 + \sigma^2)s_j + \eta]. \quad (17)$$

This criterion and the criterion $\Delta_j > 0$ yield interesting insights into the roles of various hyperparameters choices.

First (Obs#6), larger predictor learning rate α_p can play an advantageous role by loosening the upper bound in Eqn. 17,

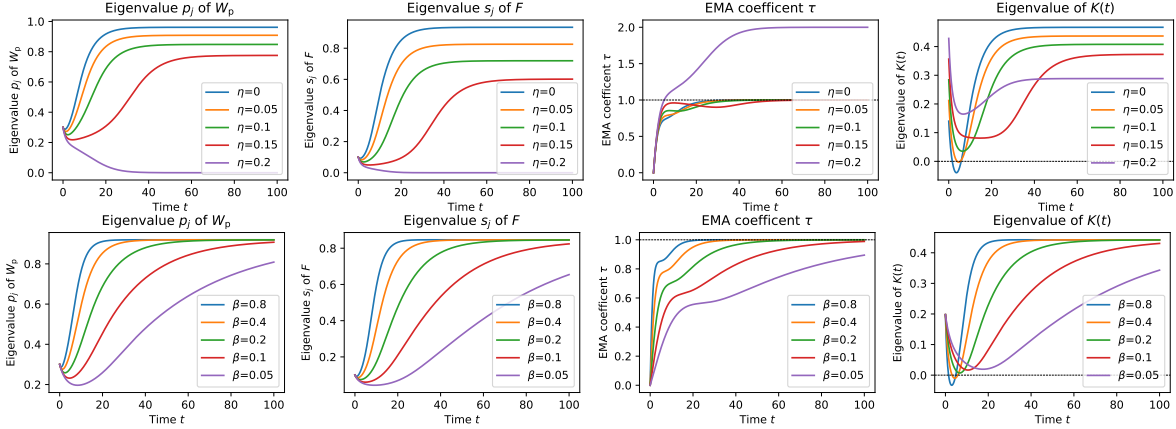


Figure 5. The role played by weight decay η and EMA β when applying symmetric regularization on W_p . The learning rate $\alpha = 0.01$. Both terms boost the eigenvalue of $K(t)$ to above 0 so that eigen space alignment could happen (Theorem 3), but also come with different trade-offs. Here $\beta = 0.4$ so that $\alpha\beta = 0.004 = 1 - \gamma_a$ where $\gamma_a = 0.996$ as in BYOL. **Top row (Weight Decay η)**: A large η boost the eigenvalue of $K(t)$ up, but substantially decreases the final converging eigenvalues p_j and s_j (i.e., the final features are not salient), or even drags them to zero (no training happens). **Bottom row (EMA β)**. A small EMA β also boost the eigenvalue of $K(t)$, but the training converges much slower. Here $\eta = 0.04$ so that $\eta\alpha$ equals to the weight decay ($\bar{\eta} = 0.0004$) in our STL-10 experiments.

making it easier to satisfy. Second (Obs#7), increasing η also has the same effect.

Role of EMA. Without EMA, $\tau \equiv 1$ and (Eqn. 17) may not hold initially when p_j is small. The reason is Δ_j is to leading order linear in p_j when $\tau = 1$ while the right hand side is to leading order $s_j \sim p_j^2$, so the left hand side has a larger contribution from p_j than the right.

EMA resolves this as follows. When the training begins, s_j is often quite small, and τ remains small since W changes rapidly. When p_j grows to the fixed point $p_{j+}^* \sim \tau/(1 + \sigma^2)$, the growth of s_j stops, making τ larger. This in turns sets a higher fixed point goal for p_j . This process continues until the feature is stabilized and $\tau = 1$ (Fig. 5 for details).

Therefore, EMA can serve as an *automatic curriculum* (Obs#8): it sets an initial small goal of $\frac{\tau}{1 + \sigma^2}$ for p_j so Δ_j need only be small and positive to both drive p_j larger and satisfy Eqn. 17. Then EMA gradually sets a higher goal for p_j by increasing τ , so that p_j and s_j can grow, while keeping the eigenspaces of W_p and F aligned.

As a trade-off, a very slow EMA schedule (β small) yields a slow training procedure (Obs#9) (See Fig. 5). Also small τ leads to larger p_{j-}^* and more eigen modes can be trapped in the collapsed basin (Obs#10).

3.3. Summarization of effects of hyperparameters

We summarize the positive and negative effects of multiple hyperparameters in Tbl. 4. We next provide additional ablations and experiments to further justify our reasoning.

Different weight decay η_p and η_s . If we set a higher weight decay for the predictor (η_p) than the online net (η_s),

	No predictor bias		With predictor bias	
	sym W_p	regular W_p	sym W_p	regular W_p
<i>Weight decay only for predictor ($\bar{\eta}_p = 0.0004$ and $\bar{\eta}_s = 0$)</i>				
EMA	71.91 ± 0.70	70.54 ± 0.93	73.67 ± 0.47	70.89 ± 0.98
no EMA	71.12 ± 0.71	71.34 ± 0.63	73.01 ± 0.37	71.70 ± 0.83
<i>No weight decay for all ($\bar{\eta}_p = \bar{\eta}_s = 0$)</i>				
EMA	71.76 ± 0.28	70.62 ± 1.05	71.86 ± 0.39	70.99 ± 1.01
no EMA	43.04 ± 2.32	71.36 ± 0.44	41.36 ± 3.33	71.37 ± 0.77

Table 5. Symmetric weight works without EMA, if we set weight decay for the predictor ($\bar{\eta}_p = 0.0004$) but not the trunk ($\bar{\eta}_s = 0$) in BYOL experiment on STL-10. Report Top-1 accuracy after 100 epochs. If there is no weight decay for *all layers*, then again symmetric weight doesn't work without EMA.




Figure 6. The effects of relative learning rate α_p . If $\alpha_p > 1$, symmetric W_p with no EMA can also work. Experiments on STL-10 and CIFAR-10 (Krizhevsky et al., 2009) (100 epochs with 5 random seeds).

then p_j grows slower than s_j and it is possible that the condition of Theorem 3 can still be satisfied without using EMA. Indeed Tbl. 5 shows this is the case.

Larger learning rate of the predictor $\alpha_p > 1$. Our analysis predicts that one way to make symmetric W_p work with no EMA is to use $\alpha_p > 1$ (i.e. Theorem 3 is more easily satisfied). Fig. 6 verifies this prediction. Moreover Table 22 in Appendix of BYOL (Grill et al., 2020) also shows that $\alpha_p > 1$ is required to get BYOL working without EMA.

	Regularization factor ϵ			
	0	0.01	0.1	0.5
$\rho = 0.3$	76.77 \pm 0.24	77.11 \pm 0.35	77.86 \pm 0.16	75.06 \pm 1.10
$\rho = 0.5$	76.65 \pm 0.20	76.76 \pm 0.33	77.56 \pm 0.25	75.22 \pm 0.81

Table 6. STL-10 Top-1 after BYOL training for 100 epochs, if we use **DirectPred** (Eqn. 18). It outperforms training W_p using gradient descent (74.51% in Tbl. 3, regular W_p with EMA). EMA is used in all experiments. No predictor bias. ρ defined in Eqn. 19.

	Initial constant c_j			
	0.1	0.05	-0.05	-0.1
freq=1	46.57 \pm 18.43	65.31 \pm 18.22	77.11 \pm 0.66	76.46 \pm 0.55
freq=2	75.01 \pm 0.48	75.10 \pm 0.35	76.83 \pm 0.52	76.31 \pm 0.27

Table 7. STL-10 Top-1 Accuracy after BYOL training for 100 epochs. With different c_j . $\rho = 0.3$ and $\epsilon = 0$. EMA is used in all experiments. No predictor bias.

4. Optimization-free Predictor W_p

A direct consequence of our theory is a new method for choosing the predictor that avoids gradient descent altogether. Instead, we estimate F (the correlation matrix of predictor inputs) and directly set W_p to be a function of this, thereby avoiding both the need to align the eigenspaces of F and W_p through optimization, and the need to initialize W_p outside the basin of collapse. As we shall see, this exceedingly simple, theory motivated method also yields better performance in practice compared to gradient-based optimization of a linear predictor.

We call our method **DirectPred** which simply estimates F , computes its eigen-decomposition $\hat{F} = \hat{U}\hat{\Lambda}_F\hat{U}^\top$, where $\hat{\Lambda}_F = \text{diag}[s_1, s_2, \dots, s_d]$, and sets W_p via

$$p_j = \sqrt{s_j} + \epsilon \max_j s_j, \quad W_p = \hat{U} \text{diag}[p_j] \hat{U}^\top. \quad (18)$$

This choice is theoretically motivated by eigenspace-alignment between W_p and F (Theorem. 3) and convergence to the invariant parabola $s_j \propto p_j^2$ in Eqn. 14 with weight decay ($\eta > 0$). Here the estimate \hat{F} can be obtained by a moving average

$$\hat{F} = \rho \hat{F} + (1 - \rho) \mathbb{E}_B[\mathbf{f}\mathbf{f}^\top] \quad (19)$$

where $\mathbb{E}_B[\cdot]$ is the expectation over a batch. We also added a regularization factor proportional to a small ϵ to boost the small eigenvalues s_j so they can learn faster. In all our experiments on real-world datasets, we use ℓ_2 -normalization so the absolute magnitude of s_j doesn't matter.

Tbl. 6 shows that directly computing W_p through **DirectPred** works *better* (76.77%) than training via gradient descent (74.51% in Tbl. 3, regular W_p with EMA). Additional regularization through ϵ yields even better performance (77.38%). Different ways to estimate F (moving average or simple average) yield only small differences.

The performance of **DirectPred** also remains good over many more training epochs (Tbl. 8). Moreover, if we allow

some gradient steps in between directly setting W_p , performance becomes even better (80.28%). This might occur because the estimated \hat{F} may not be accurate enough and SGD can help correct it. This also mitigates the computational cost of eigen-decomposition.

The constant c_j . What happens if $p_j = \sqrt{\max(s_j - c_j, 0)}$ with $c_j \neq 0$? If c_j is small negative, performance is still fine but a positive c_j leads to very poor performance (Tbl. 7), likely due to many small eigen-values s_j becoming zero and therefore trapped in the collapsed basin.

Feature-dependent W_p . Note one of the advantages of using two layer predictors is that W_p can depend on the input features. We explored this idea by using a few random partitions of the input space, and within each random partition we estimated a different correlation matrix \hat{F} . The final \hat{F} is the sum of all the correlation matrices. With 6 random partitions, **DirectPred** achieves 78.20 ± 0.16 Top-1 accuracy after 100 epochs, closing performance gap to two-layer predictors (78.85% in Tbl. 3). We leave a thorough analysis of the two layer setting to future work.

ImageNet experiments. We conducted additional experiments on ImageNet (Deng et al., 2009), with our BYOL (Grill et al., 2020) implementation. We used ResNet-50 (He et al., 2016) as the backbone that produces features, followed by a projection MLP and a predictor MLP (i.e. the BYOL design). The results are summarized in Tbl. 9. As a baseline, the default 2-layer predictor (with BatchNorm and ReLU, 4096 hidden dimension, 256 input/output dimension) achieves 64.7% top-1 accuracy, and 85.8% top-5 accuracy when pre-trained for 60 epochs and then evaluated with a linear classifier on the ResNet pool-5 features. On the other hand, we find **DirectPred** can match this performance (64.4% top-1, and 85.8% top-5) *without* any gradient-based training by instead directly setting the linear predictor weights every mini-batch. For a fair comparison, we also run BYOL with a learned linear predictor. We find the performance drops to 59.4, and 82.3 respectively. These experiments demonstrate the success of **DirectPred** on STL-10 and CIFAR can also scale to ImageNet.

Summary. Thus remarkably, our theoretical analysis of non-contrastive SSL, primarily centered around a 3 dimensional nonlinear dynamical system, not only yields conceptual insights into the functional roles of complex ingredients like EMA, stop-gradients, predictors, predictor symmetry, diverse learning rates, weight decay and all their interactions, but also predicts the performance patterns of many ablation studies as well as suggests an exceedingly simple **DirectPred** method that rivals the performance of more complex predictor dynamics in real-world settings.

	Number of epochs		
	100	300	500
<i>STL-10</i>			
DirectPred	77.86 ± 0.16	78.77 ± 0.97	78.86 ± 1.15
DirectPred (freq=5)	77.54 ± 0.11	79.90 ± 0.66	80.28 ± 0.62
SGD baseline	75.06 ± 0.52	75.25 ± 0.74	75.25 ± 0.74
<i>CIFAR-10</i>			
DirectPred	85.21 ± 0.23	88.88 ± 0.15	89.52 ± 0.04
DirectPred (freq=5)	84.93 ± 0.29	88.83 ± 0.10	89.56 ± 0.13
SGD baseline	84.49 ± 0.20	88.57 ± 0.15	89.33 ± 0.27

Table 8. STL-10/CIFAR-10 Top-1 accuracy of **DirectPred**, after training for longer epochs. $\rho = 0.3$, $\epsilon = 0.1$ with EMA.

BYOL variants	Accuracy	
	Top-1	Top-5
2-layer predictor (default)	64.7	85.8
linear predictor	59.4	82.3
DirectPred	64.4	85.8

Table 9. ImageNet experiments (60 epochs) comparing **DirectPred** with BYOL (Grill et al., 2020). Without gradient-based training, **DirectPred** is able to match the performance of the default 2-layer predictor introduced by BYOL, and significantly outperform the linear predictor by 5%.

Acknowledgement

We thank Lantao Yu for helpful discussions.

References

Arora, S., Cohen, N., and Hazan, E. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, pp. 244–253. PMLR, 2018.

Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. February 2019.

Bachman, P., Hjelm, R. D., and Buchwalter, W. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019.

Bartlett, P., Helmbold, D., and Long, P. Gradient descent with identity initialization efficiently learns positive definite linear transformations by deep residual networks. In *International conference on machine learning*, pp. 521–530. PMLR, 2018.

Bromley, J., Guyon, I., LeCun, Y., Säcker, E., and Shah, R. Signature verification using a “siamese” time delay neural network. *Adv. Neural Inf. Process. Syst.*, pp. 737–737, 1994.

Brutzkus, A. and Globerson, A. Globally optimal gradient descent for a convnet with gaussian inputs. In *International conference on machine learning*, pp. 605–614. PMLR, 2017.

Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 2020.

Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020a.

Chen, X. and He, K. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020.

Chen, X., Fan, H., Girshick, R., and He, K. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.

Coates, A., Ng, A., and Lee, H. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223, 2011.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

Du, S. and Hu, W. Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning*, pp. 1655–1664. PMLR, 2019.

Du, S. S., Hu, W., and Lee, J. D. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. *arXiv preprint arXiv:1806.00900*, 2018.

Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. *ICML*, 2019.

Fetterman, A. and Albrecht, J. Understanding self-supervised and contrastive learning with “bootstrap your own latent” (byol), 2020. URL <https://untitled-ai.github.io/understanding-self-supervised-contrastive-learning.html#fn:ssup>.

Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation

- learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Kawaguchi, K. Deep learning without poor local minima. *NeurIPS*, 2016.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Lampinen, A. K. and Ganguli, S. An analytic theory of generalization dynamics and transfer learning in deep linear networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- Laurent, T. and Brecht, J. Deep linear networks with arbitrary loss: All local minima are global. In *International conference on machine learning*, pp. 2902–2907. PMLR, 2018.
- Lee, J. D., Lei, Q., Saunshi, N., and Zhuo, J. Predicting what you already know helps: Provable self-supervised learning. *arXiv preprint arXiv:2008.01064*, 2020.
- Misra, I. and Maaten, L. v. d. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6707–6717, 2020.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Pennington, J., Schoenholz, S., and Ganguli, S. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In *Advances in Neural Information Processing Systems*. 2017.
- Pennington, J., Schoenholz, S. S., and Ganguli, S. The emergence of spectral universality in deep networks. In *Artificial Intelligence and Statistics (AISTATS)*, 2018.
- Safran, I. and Shamir, O. Spurious local minima are common in two-layer relu neural networks. In *International Conference on Machine Learning*, pp. 4433–4441. PMLR, 2018.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. A mathematical theory of semantic development in deep neural networks. *Proc. Natl. Acad. Sci. U. S. A.*, May 2019.
- Tian, Y. An analytical formula of population gradient for two-layered relu network and its applications in convergence and critical point analysis. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*, pp. 3404–3413. JMLR. org, 2017.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive learning, multi-view redundancy, and linear models. *arXiv preprint arXiv:2008.10150*, 2020.

A. Section 2

Lemma 1 (Dynamics of BYOL/SimSiam). *For objective ($\mathbf{f}_1 = W\mathbf{x}_1$ and $\mathbf{f}_{2a} = W_a\mathbf{x}_2$ where W_a is EMA weight):*

$$J(W, W_p) := \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p(\cdot), \mathbf{x}_1, \mathbf{x}_2 \sim p_{\text{aug}}(\cdot|\mathbf{x})} [\|W_p \mathbf{f}_1 - \text{StopGrad}(\mathbf{f}_{2a})\|_2^2] \quad (20)$$

Let $X = \mathbb{E}[\bar{\mathbf{x}}\bar{\mathbf{x}}^\top]$ where $\bar{\mathbf{x}}(\mathbf{x}) := \mathbb{E}_{\mathbf{x}' \sim p_{\text{aug}}(\cdot|\mathbf{x})}[\mathbf{x}']$ is the average augmented view of a data point \mathbf{x} and $X' = \mathbb{E}_{\mathbf{x}}[\mathbb{V}_{\mathbf{x}'|\mathbf{x}}[\mathbf{x}']]$ is the covariance matrix $\mathbb{V}_{\mathbf{x}'|\mathbf{x}}[\mathbf{x}']$ of augmented views \mathbf{x}' conditioned on \mathbf{x} , subsequently averaged over the data \mathbf{x} . The dynamics is the following:

$$\dot{W}_p = -\frac{\partial J}{\partial W_p} = -W_p W (X + X') W^\top + W_a X W^\top \quad (21)$$

$$\dot{W} = -\frac{\partial J}{\partial W} = -W_p^\top W_p W (X + X') + W_p^\top W_a X \quad (22)$$

Proof. Note that

$$(W_p \mathbf{f}_1 - \mathbf{f}_{2a})^\top (W_p \mathbf{f}_1 - \mathbf{f}_{2a}) \quad (23)$$

$$= \mathbf{f}_1^\top W_p^\top W_p \mathbf{f}_1 - \mathbf{f}_{2a}^\top W_p \mathbf{f}_1 - \mathbf{f}_1^\top W_p^\top \mathbf{f}_{2a} + \mathbf{f}_{2a}^\top \mathbf{f}_{2a} \quad (24)$$

$$= \text{tr}(W_p^\top W_p \mathbf{f}_1 \mathbf{f}_1^\top) - \text{tr}(W_p \mathbf{f}_1 \mathbf{f}_{2a}^\top) - \text{tr}(W_p^\top \mathbf{f}_{2a} \mathbf{f}_1^\top) + \text{tr}(\mathbf{f}_{2a} \mathbf{f}_{2a}^\top) \quad (25)$$

Let $F_1 = \mathbb{E}[\mathbf{f}_1 \mathbf{f}_1^\top] = W(X + X')W^\top$ where $X = \mathbb{E}_{\mathbf{x}}[\bar{\mathbf{x}}\bar{\mathbf{x}}^\top]$ and $X' = \mathbb{E}_{\mathbf{x}}[\mathbb{V}_{\mathbf{x}'|\mathbf{x}}[\mathbf{x}']]$, $F_{1,2a} = \mathbb{E}[\mathbf{f}_1 \mathbf{f}_{2a}^\top]$, $F_{2a,1} = \mathbb{E}[\mathbf{f}_{2a} \mathbf{f}_1^\top] = F_{1,2a}^\top$ and $F_{2a} = \mathbb{E}[\mathbf{f}_{2a} \mathbf{f}_{2a}^\top]$. This leads to:

$$J(W, W_p) = \frac{1}{2} [\text{tr}(W_p^\top W_p F_1) - \text{tr}(W_p F_{1,2a}) - \text{tr}(F_{1,2a} W_p) + \text{tr}(F_{2a})] \quad (26)$$

Taking partial derivative with respect to W_p and we get the gradient update rule:

$$\dot{W}_p = -\frac{\partial J}{\partial W_p} = -W_p F_1 + F_{1,2a}^\top \quad (27)$$

Now we take the derivative with respect to W . Note that we have stop-gradient in \mathbf{f}_{2a} , so we would like to be careful when taking derivatives. We first compute $\partial J / \partial F_1$ and $\partial J / \partial F_{1,2a}$. Note that both F_1 and $F_{1,2a}$ contains W , due to the fact that we have stop gradient, F_1 is a quadratic form of W but $F_{1,2a}$ is a *linear* form of W . This is critical.

$$\frac{\partial J}{\partial F_1} = \frac{1}{2} W_p^\top W_p \quad (28)$$

$$\frac{\partial J}{\partial F_{1,2a}} = -W_p^\top \quad (29)$$

Let $W = [w_{ij}]$ and $X = \mathbb{E}[\bar{\mathbf{x}}\bar{\mathbf{x}}^\top]$ (X_{tot} and X' are defined similarly). We have $F_1 = W(X + X')W^\top$ and $F_{1,2a} = W X W_a^\top$. So we have:

$$\frac{\partial J}{\partial w_{ij}} = \sum_{kl} \left[\frac{\partial J}{\partial F_1} \right]_{kl} \frac{\partial [F_1]_{kl}}{\partial w_{ij}} + \sum_{kl} \left[\frac{\partial J}{\partial F_{1,2a}} \right]_{kl} \frac{\partial [F_{1,2a}]_{kl}}{\partial w_{ij}} \quad (30)$$

Let $C = X + X'$, here we have:

$$\sum_{kl} \left[\frac{\partial J}{\partial F_1} \right]_{kl} \frac{\partial [F_1]_{kl}}{\partial w_{ij}} = \sum_{kl} \left[\frac{\partial J}{\partial F_1} \right]_{kl} \sum_{mn} \frac{\partial w_{km} c_{mn} w_{ln}}{\partial w_{ij}} \quad (31)$$

$$= \sum_{kl} \left[\frac{\partial J}{\partial F_1} \right]_{kl} \left(\delta(i=k) \sum_n c_{jn} w_{ln} + \delta(i=l) \sum_m w_{km} c_{mj} \right) \quad (32)$$

$$= \sum_l \left[\frac{\partial J}{\partial F_1} \right]_{il} \sum_n c_{jn} w_{ln} + \sum_k \left[\frac{\partial J}{\partial F_1} \right]_{ki} \sum_m w_{km} c_{mj} \quad (33)$$

$$= \left[\frac{\partial J}{\partial F_1} W C^\top + \frac{\partial J}{\partial F_1^\top} W C \right]_{ij} \quad (34)$$

Similarly (note that we don't take derivative with respect to W_a):

$$\sum_{kl} \left[\frac{\partial J}{\partial F_{1,2a}} \right]_{kl} \frac{\partial [F_{1,2a}]_{kl}}{\partial w_{ij}} = \left[\frac{\partial J}{\partial F_{1,2a}} W_a X^\top \right]_{ij} \quad (35)$$

So we have:

$$\dot{W} = -\frac{\partial J}{\partial W} = -W_p^\top W_p W (X + X') + W_p^\top W_a X \quad (36)$$

After some manipulation, we finally arrive at the following gradient update rule:

$$\dot{W}_p = [-W_p W (X + X') + W_a X] W^\top - \eta W_p \quad (37)$$

$$\dot{W} = W_p^\top [-W_p W (X + X') + W_a X] - \eta W \quad (38)$$

□

Remarks. For symmetric loss:

$$J(W, W_p) := \frac{1}{4} \mathbb{E}_{\mathbf{x} \sim p(\cdot), \mathbf{x}_1, \mathbf{x}_2 \sim p_{\text{aug}}(\cdot|\mathbf{x})} [\|W_p \mathbf{f}_1 - \text{StopGrad}(\mathbf{f}_{2a})\|_2^2 + \|W_p \mathbf{f}_2 - \text{StopGrad}(\mathbf{f}_{1a})\|_2^2] \quad (39)$$

The update rule is done by swapping subscript 1 and 2 in the update rule of W_p (here $F_2 = \mathbb{E}[\mathbf{f}_2 \mathbf{f}_2^\top]$):

$$\dot{W}_p = -\frac{\partial J}{\partial W_p} = -\frac{1}{2} W_p (F_1 + F_2) + \frac{1}{2} (F_{2a,1} + F_{1a,2}) \quad (40)$$

Under the large batch limit, it is the same as Eqn. 37.

Theorem 1 (Invariance of the Gradient Update). *The gradient update rules (Eqn. 2 and Eqn. 3) has the following invariance (where the symmetric matrix C depends on initialization):*

$$W(t) W^\top(t) = W_p^\top(t) W_p(t) + e^{-2\eta t} C \quad (41)$$

Proof. From Eqn. 38 and Eqn. 37, we know that

$$\alpha_p^{-1} W_p^\top \dot{W}_p + \alpha_p^{-1} \eta W_p^\top W_p = \dot{W} W^\top + \eta W W^\top \quad (42)$$

Taking transpose and we have:

$$\alpha_p^{-1} \dot{W}_p^\top W_p + \alpha_p^{-1} \eta W_p^\top W_p = W \dot{W}^\top + \eta W W^\top \quad (43)$$

Adding them together and multiply both side with $e^{2\eta t}$:

$$\alpha_p^{-1} \frac{d}{dt} (e^{2\eta t} W_p^\top W_p) = \frac{d}{dt} (e^{2\eta t} W W^\top) \quad (44)$$

This leads to $\alpha_p^{-1} e^{2\eta t} W W^\top = e^{2\eta t} W_p^\top W_p + C$, or $W W^\top = \alpha_p^{-1} W_p^\top W_p + e^{-2\eta t} C$. □

Theorem 2 (No-stop gradient will not work). *Without no-stop gradient and $W_a = W$, the gradient update rule for W is the following (here $\tilde{W}_p = W_p - I$):*

$$\frac{d}{dt} \text{vec}(W) = -(X' \otimes W_p^\top W_p + X \otimes \tilde{W}_p^\top \tilde{W}_p) \text{vec}(W) \quad (45)$$

which leads to $W(t) \rightarrow 0$ when X' and W_p are full-rank.

Proof. Note that if we don't have stop gradient and $W_a = W$, then we have additional terms (and we also need to compute $\partial J/\partial F_2$). Let $\tilde{W}_p = W_p - I$ and we have:

$$\dot{W} = -\frac{\partial J}{\partial W} = -W_p^\top W_p W(X + X') + (W_p^\top + W_p)WX - W(X + X') \quad (46)$$

$$= -(W_p^\top W_p + I)WX' - (W_p^\top W_p - W_p^\top - W_p + I)WX \quad (47)$$

$$= -(W_p^\top W_p + I)WX' - (W_p - I)^\top (W_p - I)WX \quad (48)$$

$$= -(W_p^\top W_p + I)WX' - \tilde{W}_p^\top \tilde{W}_p WX \quad (49)$$

With $\text{vec}(AXB) = (B^\top \otimes A)\text{vec}(X)$ and we see:

$$\frac{d}{dt}\text{vec}(W) = -\left[X' \otimes (W_p^\top W_p + I) + X \otimes \tilde{W}_p^\top \tilde{W}_p\right]\text{vec}(W) \quad (50)$$

Note that the leading terms are always negative semi-definite if X' is full-rank, and there is no chance for W to learn any meaningful features. \square

B. Section 3

Isometric assumptions. Now we use the assumption that $X = I$ and $X' = \sigma^2 I$, which leads to

$$\dot{F} = \dot{W}XW^\top + WX\dot{W}^\top = -(1 + \sigma^2)(W_p^\top W_p F + FW_p^\top W_p) + W_p^\top W_a W^\top + WW_a^\top W_p \quad (51)$$

here $F = WXW^\top = WW^\top$. If we also have weight decay $-\eta W$ for W , then we have:

$$\dot{F} = -(1 + \sigma^2)(W_p^\top W_p F + FW_p^\top W_p) + W_p^\top W_a W^\top + WW_a^\top W_p - 2\eta F \quad (52)$$

or using anticommutator $\{A, B\} := AB + BA$:

$$\dot{F} = -(1 + \sigma^2)\{F, W_p^\top W_p\} + W_p^\top W_a W^\top + WW_a^\top W_p - 2\eta F \quad (53)$$

Similarly, for W_p we have:

$$\dot{W}_p = -\alpha_p(1 + \sigma^2)W_p F + \alpha_p \tau F - \eta W_p \quad (54)$$

EMA assumption (Assumption 1). Now we further study the effect of EMA. To model it, we just let $W_a = \tau W$ where $\tau < 1$ is a coefficient that measure how much EMA attenuates W . If $\tau = 1$ then $W_a = W$ and there is no EMA. Note that τ is not the same as the EMA parameter $1 - \gamma_a$, which is often set to be a fixed 0.004 (or $1 - 0.996$). Instead, $\tau = \tau(t)$ is a changing parameter depends on how quickly $W = W(t)$ grows over time. If W remains stable, then $\tau \approx 1$; if W grows rapidly, then τ becomes small.

Fig. 7 shows that this assumption is largely correct.

Under this condition, using $F = WXW^\top = WW^\top$, the dynamics becomes (Now we also put weight decay for W_p):

$$\dot{W}_p = -\alpha_p(1 + \sigma^2)W_p F + \alpha_p \tau F - \eta W_p \quad (55)$$

$$\dot{F} = -(1 + \sigma^2)(W_p^\top W_p F + FW_p^\top W_p) + \tau(W_p^\top F + FW_p) - 2\eta F \quad (56)$$

Derivation of Fixed point of Eqn. 2. Given the dynamics Eqn. 55 we now want to check its fixed point:

$$-\alpha_p(1 + \sigma^2)W_p F + \alpha_p \tau F - \eta W_p = 0 \quad (57)$$

for some PSD matrix F . For convenience, let $\eta' = \eta/\alpha_p$. Since F is always PSD, we have eigendecomposition $F = U\Lambda U^\top$. Left-multiplying U and right-multiplying U^\top , we have:

$$(1 + \sigma^2)\bar{W}_p \Lambda + \eta' \bar{W}_p = \tau \Lambda \quad (58)$$

where $\bar{W}_p := U^\top W_p U$. Let $\Lambda' = (1 + \sigma^2)\Lambda + \eta' I$ is a diagonal matrix with all positive diagonal element since $\eta' > 0$. Therefore, we have:

$$\bar{W}_p \Lambda' = \tau \Lambda \quad (59)$$




Figure 7. Check the validity of EMA assumption (Assumption 1) with different EMA coefficients γ_a for BYOL dynamics with $X = I$ and $X' = \sigma^2 I$ (Assumption 2). $\sigma = 0.03$. All experiments are run 10 times to get mean and standard derivation (shaded area). We could see the EMA assumption is largely correct. Even at the region with γ_a close to 1 (e.g., 0.996) and large η , the normalized correlation between W_a and W are still high (~ 0.9). Note that throughout our analysis, the initial value of $W_a(0) = 0$. **Left:** weight decay $\eta = 0$, **Middle:** $\eta = 0.01$, **Right:** $\eta = 0.1$.

and thus $\bar{W}_p = \tau \Lambda (\Lambda')^{-1}$ is a symmetric matrix and so does $W_p = U \bar{W}_p U^\top$. When $\eta = 0$ and F has zero eigenvalues, W_p can have infinite solutions (or fixed points), and some of them might not be symmetric.

Symmetrization of W_p . Now we need to assume W_p is symmetric and also symmetrize its dynamics, which yields (here $\{A, B\} := AB + BA$):

$$\begin{aligned} \dot{W}_p &= -\frac{\alpha_p}{2}(1 + \sigma^2)\{W_p, F\} + \alpha_p \tau F - \eta W_p \\ \dot{F} &= -(1 + \sigma^2)\{W_p^2, F\} + \tau\{W_p, F\} - 2\eta F \end{aligned} \quad (60)$$

Note that the asymmetric dynamic might be interesting and we will leave it later.

B.1. Section 3.1

Theorem 3 (Alignment of Eigenspace). *Under the dynamics of Eqn. 60, the commutator $[F, W_p] := FW_p - W_p F$ satisfies:*

$$\frac{d}{dt}[F, W_p] = -[F, W_p]K - K[F, W_p] \quad (61)$$

where

$$K = K(t) = (1 + \sigma^2) \left[\frac{\alpha_p}{2} F(t) + W_p^2(t) - \frac{\tau}{1 + \sigma^2} W_p(t) \right] + \frac{3}{2} \eta I \quad (62)$$

If $\max_{t \geq 0} \lambda_{\min}[K(t)] = \lambda_0 > 0$, then the commutator $\|[F(t), W_p(t)]\|_F \leq e^{-2\lambda_0 t} \|[F(0), W_p(0)]\|_F \rightarrow 0$, i.e., the eigenspace of W_p gradually aligns with F .

Proof. Let's compute the commutator $L := [F, W_p] := FW_p - W_p F$ and its time derivative. First we have:

$$F\dot{W}_p - \dot{W}_p F = -\frac{\alpha_p}{2}(1 + \sigma^2)(FL + LF) - \eta L \quad (63)$$

Then we have

$$\dot{F}W_p - W_p\dot{F} = -(1 + \sigma^2)(W_p^2 L + L W_p^2) + \tau(W_p L + L W_p) - 2\eta L \quad (64)$$

So we have

$$\dot{L} = F\dot{W}_p + \dot{F}W_p - (W_p\dot{F} + \dot{W}_p F) = -KL - LK \quad (65)$$

where

$$K = K(t) = (1 + \sigma^2) \left[\frac{\alpha_p}{2} F + W_p^2 - \frac{\tau}{1 + \sigma^2} W_p \right] + \frac{3}{2} \eta I \quad (66)$$

is a symmetric matrix. We can write the close form solution for L :

$$\text{vec}(L(t)) = A(t)\text{vec}(L(0)) \quad (67)$$

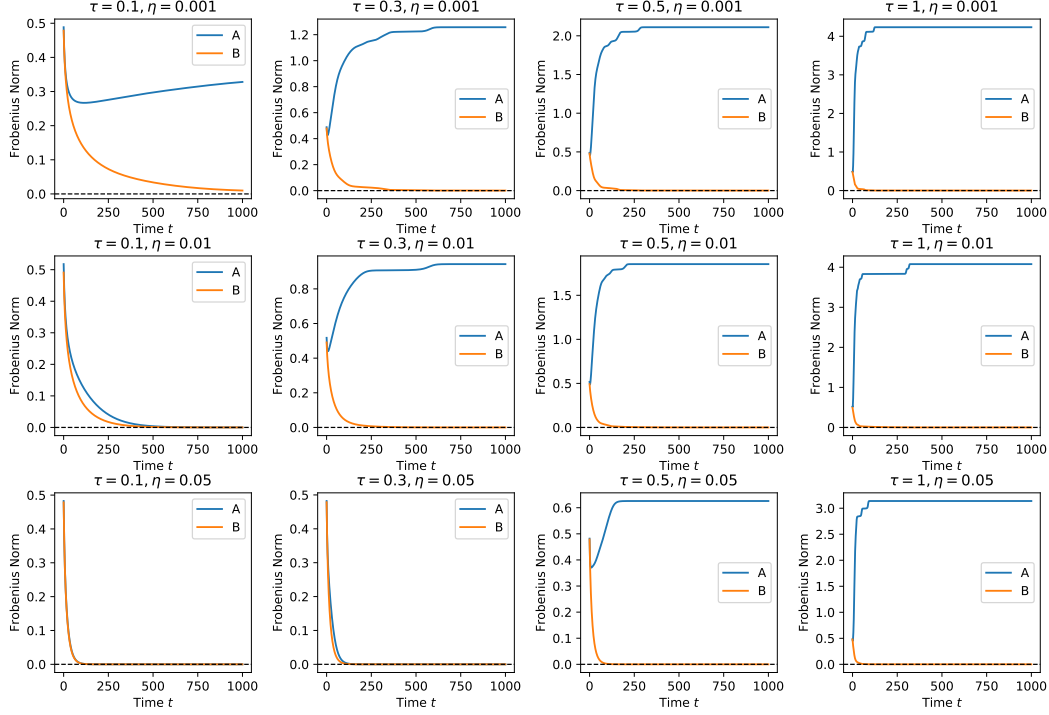


Figure 8. Dynamics of the symmetric $A := (W_p + W_p^\top)/2$ and asymmetric part $B := (W_p - W_p^\top)/2$ of W_p under different *time-independent* τ of Eqn. 55. Each row is a different weight decay η (i.e., $\eta = 0.001, 0.01$ and 0.05). When η is large and/or τ is small, $\|A\|_F$ can also be dragged to zero, which is consistent with analysis in Sec. 3.2 (Obs#4 and Obs#5). On the other hand, $\|B\|_F$ always seems to vanish over time. In this numerical simulation, we set $F = W_p^\top W_p$ following invariant in Theorem 1 with $C = 0$.

where $A(t) = \exp\left(-\int_0^t K(t') \oplus K(t') dt'\right)$ and $K(t') \oplus K(t') := I \otimes K(t') + K(t') \otimes I$ is the Kronecker sum.

If $\lambda_{\min}(K(t)) \geq \lambda_0 > 0$ for all t , then $\lambda_{\min}[K(t') \oplus K(t')] \geq 2\lambda_0$. For any unit vector \mathbf{v} , we have:

$$\mathbf{v}^\top \left(\int_0^t K(t') \oplus K(t') dt' \right) \mathbf{v} = \int_0^t \mathbf{v}^\top K(t') \oplus K(t') \mathbf{v} dt' \geq \int_0^t 2\lambda_0 \mathbf{v}^\top \mathbf{v} dt = 2\lambda_0 t \quad (68)$$

Therefore, $\lambda_{\max}[A(t)] \leq e^{-2\lambda_0 t}$ and $A(t)$ is a contraction and

$$\|L\|_F = \|\text{vec}^\top(L)\|_2 \leq \lambda_{\max}[A(t)] \|\text{vec}(L(0))\|_2 \leq e^{-2\lambda_0 t} \|L(0)\|_F \rightarrow 0 \quad (69)$$

This means that W_p and F can commute, and the eigen space of W_p and F will gradually align. \square

Remark. Fig. 9 shows numerical simulation of the symmetrized dynamics (Eqn. 60). If $K(t)$ has negative eigenvalues, then even if W_p and F have already approximately aligned, the dynamics is also unstable and might diverge due to noise and/or numerical instability.

Fig. 8 shows a numerical simulation of Eqn. 55 (dynamics with Assumption 1 and Assumption 2 but without the symmetric dynamics). We can clearly see that the asymmetric component converges to zero.

When eigenspace aligns exactly. Let U be the common eigenvectors. $W_p = U\Lambda_{W_p}U^\top$ where $\Lambda_{W_p} = \text{diag}[p_1, p_2, \dots, p_d]$, $F = U\Lambda_F U^\top$ where $\Lambda_F = \text{diag}[s_1, s_2, \dots, s_d]$.

In this case, the time derivatives \dot{W}_p and \dot{F} can all be written as decoupled form: $\dot{W}_p = UG_1U^\top$ and $\dot{F} = UG_2U^\top$ where G_1 and G_2 are diagonal matrices. In other words, they are both *decoupled* into each eigen mode, and so does the future value of W_p and F . Then U won't change over time.

To see why, we consider the general case where we have a symmetric matrix $M(t)$ with eigen decomposition $M(t) = U(t)D(t)U^\top(t)$. M follows $\dot{M} = U(t)G(t)U^\top(t)$ where $G(t)$ is an arbitrary diagonal matrix.

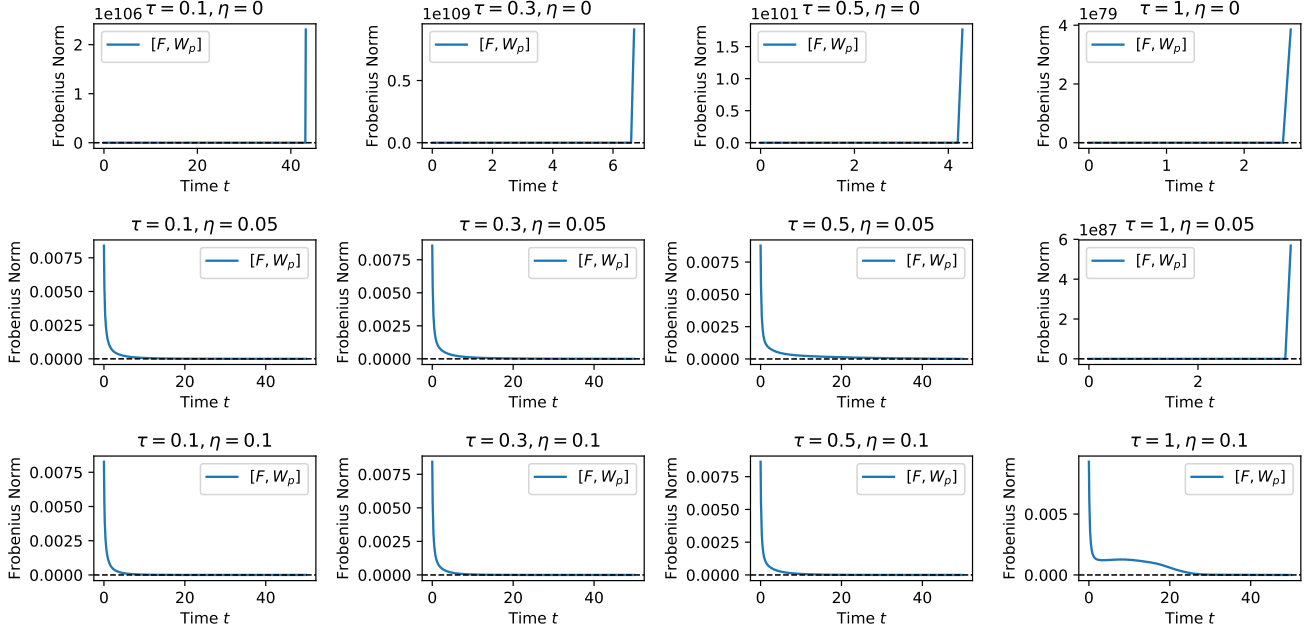


Figure 9. The norm of the communicator $[F, W_p]$ over time under different hyper-parameters (different *time-independent* τ and different weight decay η) in symmetrized dynamics Eqn. 60. When weight decay is small or zero, and/or τ is large, the norm of the communicator $\|[F, W_p]\|_F$ can shoot up (no eigenspace alignment).

To see why $\dot{U} = 0$, at each time step we have:

$$\dot{M} = \dot{U}DU^\top + U\dot{D}U^\top + U\dot{U}^\top = UGU^\top \quad (70)$$

since U is unitary, we have:

$$U^\top \dot{U}D + D\dot{U}^\top U = G - \dot{D} \quad (71)$$

Since $U^\top(t)U(t) = I$, we have $\dot{U}^\top U + U^\top \dot{U} = 0$ so $Q := U^\top \dot{U}$ is a skew-symmetric matrix and we have

$$QD - DQ = G - \dot{D} \quad (72)$$

Since the right hand side is a diagonal matrix, checking each entry and we have $q_{ij}d_j - q_{ij}d_i = 0$ for $i \neq j$. If M has distinctive eigenvalues, then we know $q_{ij} = 0$ for $i \neq j$. Q is skew-symmetric so $q_{ii} = 0$. So $Q = U^\top \dot{U} = 0$ and thus $\dot{U} = 0$. If M has duplicated eigenvalues, then we can show $q_{ij} = 0$ for any $d_i \neq d_j$. Within high-dimensional eigenspace for duplicated eigenvalues, its eigen-decomposition is not unique and we can always pick the eigenspace within each duplicated eigenspace so that $\dot{U} = 0$.

Therefore, we just multiply U^\top and U to Eqn. 60 and the system becomes decoupled. Then after some algebraic manipulation, we arrive at the following:

$$\dot{p}_j = \alpha_p(1 + \sigma^2)s_j \left[\frac{\tau}{1 + \sigma^2} - p_j \right] - \eta p_j \quad (73)$$

$$\dot{s}_j = 2(1 + \sigma^2)p_j s_j \left[\frac{\tau}{1 + \sigma^2} - p_j \right] - 2\eta s_j \quad (74)$$

Multiply Eqn. 73 with $2\alpha_p^{-1}p_j$ and subtract with Eqn. 74, we get:

$$2\alpha_p^{-1}p_j\dot{p}_j - \dot{s}_j = -2\eta\alpha_p^{-1}p_j^2 + 2\eta s_j \quad (75)$$

which gives

$$\alpha_p^{-1} \left(\frac{dp_j^2}{dt} + 2\eta p_j^2 \right) = \dot{s}_j + 2\eta s_j \quad (76)$$

$$\alpha_p^{-1} \frac{d}{dt} (e^{2\eta t} p_j^2) = \frac{d}{dt} (e^{2\eta t} s_j) \quad (77)$$

$$\alpha_p^{-1} e^{2\eta t} p_j^2 = e^{2\eta t} s_j - c_j \quad (78)$$

$$\alpha_p^{-1} p_j^2(t) = s_j(t) - e^{-2\eta t} c_j \quad (79)$$

Therefore, we have integral $s_j(t) = \alpha_p^{-1} p_j^2(t) + c_j e^{-2\eta t}$. For finite weight decay ($\eta > 0$), we could simply expect $s_j(t) \approx \alpha_p^{-1} p_j^2(t)$.

On the other hand, the dynamics of τ is:

$$\dot{W}_a = \beta(W - W_a) \quad (80)$$

Applying our assumption about EMA (Assumption 1) $W_a(t) = \tau(t)W(t)$, then we have:

$$\dot{\tau}W + \tau\dot{W} = \beta(1 - \tau)W \quad (81)$$

$$\dot{\tau}WW^\top + \tau\dot{W}W^\top = \beta(1 - \tau)WW^\top \quad (82)$$

$$2\dot{\tau}F + \tau\dot{F} = 2\beta(1 - \tau)F \quad (83)$$

When F and W_p aligns, we have \dot{F} all in the same eigen space.

$$\dot{F} = -(1 + \sigma^2)\{W_p^2, F\} + \tau\{W_p, F\} - 2\eta F \quad (84)$$

So the eigenvectors U won't change and thus we have:

$$2\dot{\tau}s_j + \tau\dot{s}_j = 2\beta(1 - \tau)s_j \quad (85)$$

or

$$\dot{\tau} = \beta(1 - \tau) - \tau \frac{\dot{s}_j}{2s_j} \quad (86)$$

which has a close form solution when $c_j = 0$. Note that in the case, we have $s_j = \alpha_p^{-1} p_j^2$ and thus $\dot{s}_j = 2\alpha_p^{-1} p_j \dot{p}_j$ and we have:

$$\dot{\tau} = \beta(1 - \tau) - \tau \frac{\dot{p}_j}{p_j} \quad (87)$$

or

$$\dot{\tau} + \tau \left(\frac{\dot{p}_j}{p_j} + \beta \right) = \beta \quad (88)$$

or

$$\frac{d}{dt} (e^{f(t)} \tau) = \beta e^{f(t)} \quad (89)$$

where $f(t) = \int (\dot{p}_j/p_j + \beta) dt = \ln p_j + \beta t$ and thus $e^{f(t)} = e^{\beta t} p_j$. Take integral on both side and we have (here $\tau(0) = 0$ is the initial condition):

$$e^{\beta t} p_j \tau = \beta \int_0^t e^{\beta t'} p_j(t') dt \quad (90)$$

which is:

$$\tau_j(t) = p_j^{-1}(t) \beta e^{-\beta t} \int_0^t p_j(t') e^{\beta t'} dt \quad (91)$$

B.2. Section 3.2

Monotonicity of p_{j-}^* with respect to η and τ . Note that

$$p_{j-}^* = \frac{\tau - \sqrt{\tau^2 - 4\eta(1 + \sigma^2)}}{2(1 + \sigma^2)} \quad (92)$$

is the (right) boundary of trivial basin $p < p_{j-}^*$ and determines the size of trivial attractive region towards $p_{j0}^* = 0$. It is dependent on η and τ . It is clear that p_{j-}^* is a increasing function of η . This means that if the weight decay η is large, so does trivial region (and more eigenvalues will be trapped to trivial solution).

On the other hand, we can compute the derivative of $g(x) = x - \sqrt{x^2 - c}$ for $c > 0$ and $x^2 > c$:

$$\frac{dg}{dx} = 1 - \frac{1}{\sqrt{1 - c/x^2}} < 0 \quad (93)$$

So $g(x)$ is a decreasing function with respect to x . Or p_{j-}^* is a decreasing function with respect to τ .

C. Section 4

Experiment setup. Unless explicitly stated, in all our experiments, we use ResNet-18 as the backbone network, two-layer MLP (with BN and ReLU) as the projector, and a linear predictor. For STL-10 and CIFAR-10, we use SGD as the optimizer with learning rate $\alpha = 0.03$, momentum 0.9, weight decay $\bar{\eta} = 0.0004$ and EMA parameter $\gamma_a = 0.996$. The batchsize is 128. Each setting is repeated 5 times to compute mean and standard derivation. We report final number as “mean \pm std”.

D. Analysis of BYOL and SimSiam learning dynamics without isotropic assumptions on data

In the main paper we focused on isotropic data assumptions to obtain analytic insights into when and why BYOL and SimSiam learning dynamics avoid representational collapse. Here we provide an alternate perspective using a different assumption, involving decoupled initial conditions, that enables us to address the case of learning with non-isotropic data. First, we recall the data generation and augmentation process. Let \mathbf{x} be a data point drawn from the data distribution $p(\mathbf{x})$ and let \mathbf{x}_1 and \mathbf{x}_2 be two augmented views of \mathbf{x} : $\mathbf{x}_1, \mathbf{x}_2 \sim p_{\text{aug}}(\cdot|\mathbf{x})$ where $p_{\text{aug}}(\cdot|\mathbf{x})$ is the augmentation distribution. Let $\Sigma^s = \mathbb{E}[\mathbf{x}_1 \mathbf{x}_1^\top]$ be the correlation matrix of a single augmented view \mathbf{x}_1 of the data \mathbf{x} , and let $\Sigma^d = \mathbb{E}[\mathbf{x}_1 \mathbf{x}_2^\top]$ be the correlation matrix between two augmented views \mathbf{x}_1 and \mathbf{x}_2 of the same data point \mathbf{x} . In the notation of the main paper, Σ^s and Σ^d can be decomposed as $\Sigma^s = X + X'$ and $\Sigma^d = X$, where $X = \mathbb{E}[\bar{\mathbf{x}} \bar{\mathbf{x}}^\top]$ and $\bar{\mathbf{x}}(\mathbf{x}) := \mathbb{E}_{\mathbf{x}' \sim p_{\text{aug}}(\cdot|\mathbf{x})}[\mathbf{x}']$ is the average augmented view of a data point \mathbf{x} . In turn $X' = \mathbb{E}_{\mathbf{x}}[\mathbb{V}_{\mathbf{x}'|\mathbf{x}}[\mathbf{x}']]$ is the covariance matrix $\mathbb{V}_{\mathbf{x}'|\mathbf{x}}[\mathbf{x}']$ of augmented views \mathbf{x}' conditioned on \mathbf{x} , subsequently averaged over the data \mathbf{x} . Intuitively, X is the correlation matrix of augmentation averaged data, while X' is the augmentation covariance matrix averaged over data.

Also recall that the BYOL learning dynamics, without weight decay, is given by

$$\dot{W} = W_p^\top (-W_p W \Sigma^s + W_a \Sigma^d) \quad (94)$$

$$\dot{W}_p = \alpha_p (-W_p W \Sigma^s + W_a \Sigma^d) W^\top \quad (95)$$

$$\dot{W}_a = \beta (-W_a + W) \quad (96)$$

SimSiam learning dynamics is a special case in which $W_a = W$ and the final equation is ignored.

We first derive exact fixed point solutions to both BYOL and SimSiam learning dynamics in this setting. We then discuss specific models for data distributions and augmentation procedures, and show how the fixed point solutions depend on both data and augmentation distributions. We then discuss how our theory reveals a fundamental role for the predictor in avoiding collapse in BYOL solutions. Finally, we derive a highly reduced three dimensional description of BYOL and SimSiam learning dynamics, assuming decoupled initial conditions, that provides considerable insights into dynamical mechanisms enabling both to avoid collapsed solutions without negative pairs to force apart representations of different objects.

D.1. The fixed point structure of BYOL and SimSiam learning dynamics.

Examining equation 94-equation 96, we find sufficient conditions for a fixed point given by $W_p W \Sigma^s = W_a \Sigma^d$ and $W = W_a$. Note these are sufficient conditions for fixed points of both BYOL and SimSiam. Inserting the second equation into the first and right multiplying both sides by $[\Sigma^s]^{-1}$ (assuming Σ^s is invertible), yields a manifold of fixed point solutions in W_1 and W_2 satisfying the nonlinear equation

$$W_p W = W \Sigma^d [\Sigma^s]^{-1}. \quad (97)$$

This constitutes a set of $n_1 \times n_2$ nonlinear equations in $(n_1 \times n_2) + (n_2 \times n_2)$ unknowns, yielding generically a nonlinear manifold of solutions in W_1 and W_2 of dimensionality $n_2 \times n_2$ corresponding to the number of predictor parameters. For concreteness, we will assume that $n_2 \leq n_1$, so that the online and target networks perform dimensionality reduction. Then a special class of solutions to equation 97 can be obtained by assuming the n_2 rows of W correspond to n_2 left-eigenvectors of $\Sigma^d [\Sigma^s]^{-1}$ and W_p is a diagonal matrix with the corresponding eigenvalues. This special class of solutions can then be generalized by a transformation $W_p \rightarrow S W_p S^{-1}$ and $W \rightarrow S W$ where S is any invertible n_2 by n_2 matrix. Indeed this transformation is a symmetry of equation 97, which defines the solution manifold. In addition to these families of solutions, the collapsed solution $W = W_p = W_a = 0$ also exists.

D.2. Illustrative models for data and data augmentation

The above section suggests that the top eigenmodes of $\Sigma^d [\Sigma^s]^{-1}$ control the non-collapsed solutions. Here we make this result more concrete by giving illustrative examples of data distributions and data augmentation procedures, and the resulting properties of $\Sigma^d [\Sigma^s]^{-1}$.

Multiplicative scrambling. Consider for example a multiplicative subspace scrambling model. In this model, data augmentation scrambles a subspace by multiplying by a random Gaussian matrix, while identically preserving the orthogonal complement of the subspace. In applications, the scrambled subspace could correspond to a space of nuisance features, while the preserved subspace could correspond to semantically important features. Indeed many augmentation procedures, including random color distortions and blurs, largely preserve important semantic information, like object identity in images.

More precisely, we consider a random scrambling operator A which only scrambles data vectors \mathbf{x} within a fixed k dimensional subspace spanned by the orthonormal columns of the $n_0 \times k$ matrix U . Within this subspace, data vectors are scrambled by a random Gaussian $k \times k$ matrix B . Thus A takes the form $A = P^c + U B U^T$ where $P^c = I - U U^T$ is a projection operator onto the $n_0 - k$ dimensional conserved, semantically important, subspace orthogonal to the span of the columns of U , and the elements of B are i.i.d. zero mean unit variance Gaussian random variables so that $\mathbb{E}[B_{ij} B_{kl}] = \delta_{ik} \delta_{jl}$. Under this simple model, the augmentation average $\bar{\mathbf{x}}(\mathbf{x}) := \mathbb{E}_{\mathbf{x}' \sim p_{\text{aug}}(\cdot|\mathbf{x})}[\mathbf{x}']$ becomes $\bar{\mathbf{x}}(\mathbf{x}) = P^c \mathbf{x}$. Thus, intuitively, under multiplicative subspace scrambling, the only aspect of a data vector that survives averaging over augmentations is the projection of this data vector onto the preserved subspace. Then the correlation matrix of two different augmented views is $\Sigma^d = P^c \Sigma^x P^c$ while the correlation matrix of two identical views is $\Sigma^s = \Sigma^x$ where $\Sigma^x \equiv \mathbb{E}_{\mathbf{x} \sim p(\cdot)}[\mathbf{x} \mathbf{x}^T]$ is the correlation matrix of the data distribution. Thus non-collapsed solutions of both BYOL and SimSiam can correspond to principal eigenmodes of $\Sigma^d [\Sigma^s]^{-1} = P^c \Sigma^x P^c [\Sigma^x]^{-1}$. In the special case in which P^c commutes with Σ^x , we have the simple result that $\Sigma^d [\Sigma^s]^{-1} = P^c$, which is completely independent of the data correlation matrix Σ^x . Thus in this simple setting BYOL and SimSiam can learn the subspace of features that are identically conserved under data augmentation, independent of how much data variance there is in the different dimensions of this conserved subspace.

Additive scrambling. We also consider, as an illustrative example, data augmentation procedures which simply add Gaussian noise with a prescribed noise covariance matrix Σ^n . Under this model, we have $\Sigma^s = \Sigma^x + \Sigma^n$ while $\Sigma^d = \Sigma^x$. Thus in this setting, BYOL learns principal eigenmodes of $\Sigma^d [\Sigma^s]^{-1} = \Sigma^x [\Sigma^x + \Sigma^n]^{-1}$. Thus intuitively, dimensions with larger noise variance are attenuated in learned BYOL representations. On the otherhand, correlations in the data that are not attenuated by noise are preferentially learned, but the degree to which they are learned is not strongly influenced by the magnitude of the data correlation (i.e. consider dimensions that lie along small eigenvalues of Σ^n). Note that in the main paper we focused on the case where $\Sigma^x = I$ and $\Sigma^n = \sigma^2 I$.

D.3. The importance of the predictor in BYOL and SimSiam.

Here we note that our theory explains why the predictor plays a crucial role in BYOL and SimSiam learning in this simple setting, as is observed empirically in more complex settings. To see this, we can model the removal of the predictor by simply setting $W_p = I$ in all the above equations. The fixed point solutions then obey $W = W\Sigma^d[\Sigma^s]^{-1}$. This will only have nontrivial, non-collapsed solutions if $\Sigma^d[\Sigma^s]^{-1}$ has eigenvectors with eigenvalue 1. Rows of W consisting of linear combinations of these eigenvectors will then constitute non-collapsed solutions.

This constraint of eigenvalue 1 yields a much more restrictive condition on data distributions and augmentation procedures for BYOL and SimSiam to have non-collapsed solutions. It can however be satisfied in multiplicative scrambling if an eigenvector of the data matrix Σ^x lies in the column space of the projection operator P^c (in which case it is an eigenvector of eigenvalue 1 of $\Sigma^d[\Sigma^s]^{-1} = P^c\Sigma^xP^c[\Sigma^x]^{-1}$). This condition cannot however be generically satisfied for additive scrambling case, in which generically all the eigenvalues of $\Sigma^d[\Sigma^s]^{-1} = \Sigma^x[\Sigma^x + \Sigma^n]^{-1}$ are less than 1. In this case, without a predictor, it can be checked that the collapsed solution $W = W_a = 0$ is stable.

Thus overall, in this simple setting, our theory provides conceptual insight into how the introduction of a predictor is crucial for creating new non-collapsed solutions for both BYOL and SimSiam, even though the predictor confers no new expressive capacity in allowing the online network to match the target network.

D.4. Reduction of BYOL learning dynamics to low dimensions

The full learning dynamics in equation 94 to equation 96 constitutes a set of high dimensional nonlinear ODEs which are difficult to solve from arbitrary initial conditions. However, there is a special class of *decoupled* initial conditions which permits additional insight. Consider the special case in which Σ^s and Σ^d commute, and so are simultaneously diagonalizable and share a common set of eigenvectors, which we denote by $\mathbf{u}^\alpha \in \mathbb{R}^{n_0}$. Consider also a special set of initial conditions where each row of W and the corresponding row of W_a are both proportional to one of the eigenmodes \mathbf{u}^α , with scalar proportionality constants w^α and w_a^α respectively, and W_p is diagonal, with the corresponding diagonal element given by w_p^α . Then it is straightforward to see that under the dynamics in equation 94 to equation 96, that the structure of this initial condition will remain the same, with only the scalars w^α , w_a^α and w_p^α changing over time. Moreover, the scalars decouple across the different indices α , and the dynamics are driven by the eigenvalues λ_s^α and λ_d^α of Σ_s and Σ_d respectively. Inserting this special class of initial conditions into the dynamics in equation 94 to equation 96, and dropping the α index, we find the dynamics of the triplet of scalars is given by

$$\frac{dw_p}{dt} = \alpha_p [w_a \lambda_d - w_p w \lambda_s] w \quad (98)$$

$$\frac{dw}{dt} = w_p [w_a \lambda_d - w_p w \lambda_s] \quad (99)$$

$$\frac{dw_a}{dt} = \beta(-w_a + w). \quad (100)$$

Alternatively, this low dimensional dynamics can be obtained from equation 94 to equation 96 not only by considering a special class of decoupled initial conditions, but also by considering the special case where every matrix is simply a 1 by 1 matrix, making the scalar replacements $W \rightarrow w$, $W_p \rightarrow w_p$, $W_a \rightarrow w_a$, $\Sigma^s \rightarrow \lambda_s$, and $\Sigma^d \rightarrow \lambda_d$. Note furthermore that this 3 dimensional dynamical system is equivalent to that studied in the main paper under the change of variables $s = w^2$ and $\tau = w_a/w$ and the special case of $\lambda_s = 1 + \sigma^2$ and $\lambda_d = 1$.

The fixed point conditions of this dynamics are given by $w_a = w$ and $w_p w = w_a \lambda_d \lambda_s^{-1}$. Thus the collapsed point $w = w_p = w_a = 0$ is a solution. Additionally $w_p = \lambda_d \lambda_s^{-1}$ and $w = w_a$ taking any value is also a family of non-collapsed solutions. We can understand the three dimensional dynamics intuitively as follows when β is much less than both 1 and α_p , so that the dynamics of w and w_p are very fast relative to the dynamics of w_a . In this case, the target network evolves very slowly compared to the online network, as is done in practice. For simplicity we use the same learning rate for the predictor as we do for the online network (i.e. $\alpha_p = 1$). In this situation, we can treat w_a as approximately constant on the fast time scale over which the online and predictor weights w and w_p evolve. Then the joint dynamics in equation 98 and equation 99 obeys gradient descent on the error function

$$E = \frac{\lambda_s}{2} (w_a \lambda_d \lambda_s^{-1} - w_p w)^2. \quad (101)$$

Iso-contours of constant error are hyperbolas in the w by w_p plane, and for fixed w_a , the origin $w = w_p = 0$ is a saddle




Figure 10. A visualization of BYOL dynamics in low dimensions. **Left:** Black arrows denote the vector field of the flow in the w and w_p plane of online and predictor weights in Eqns. 98 and 99 when the target network weight w_a is fixed to 1. For all 3 panels, $\lambda_s = 1$, $\lambda_d = 1/2$, and $\alpha_p = \beta = 1$. All flow field vectors are normalized to unit length to indicate direction of flow alone. The red curve shows the hyperbolic manifold of stable fixed points $w_p w = w_a \lambda_d \lambda_s^{-1}$, while the red point at the origin is an unstable fixed point. For a fixed target network, the online and predictor weights will cooperatively amplify each other to escape the collapsed solution at the origin. **Middle:** A visualization of the full low dimensional BYOL dynamics in Eqns 98-100 when the online and predictor weights are tied so that $w = w_p$. The green curve shows the nullcline $w_a = w$ corresponding to $\frac{dw_a}{dt} = 0$ and the blue curve shows part of the nullcline $\frac{dw}{dt} = 0$ corresponding to $w^2 = w_a \lambda_d \lambda_s^{-1}$. The intersection of these two nullclines yields two fixed points (red dots): an unstable collapsed solution at the origin $w = w_a = 0$, and a stable non-collapsed solution with $w_a = w$ and $w = \lambda_d \lambda_s^{-1}$. **Right:** A visualization of dynamics in Eqns 98-100 when the predictor is removed, so that w_2 is fixed to 1. The resulting two dimensional flow field on w and w_a is shown (black arrows). The green curve shows the nullcline $w = w_a$ corresponding to $\frac{dw_a}{dt} = 0$, while the blue curve shows the nullcline $w = w_a \lambda_d \lambda_s^{-1}$. The slope of this nullcline is $\lambda_s \lambda_d^{-1} > 1$. The resulting nullcline structure yields a single fixed point at the origin which is stable. Thus there only exists a collapsed solution. In the special case where $\lambda_s \lambda_d^{-1} = 1$, the two nullclines coincide, yielding a one dimensional manifold of solutions.

point, yielding an unstable fixed point (see Fig. 10 (left)). From generic initial conditions, w and w_p will then cooperatively amplify each other to rapidly escape the collapsed solution at the origin, and approach the zero error hyperbolic contour $w_p w = w_a \lambda_d \lambda_s^{-1}$ where w_a is close to its initial value. Then the slower target network w_a will adjust, slowly moving this contour until $w_a = w$. The more rapid dynamics of w and w_p will hug the moving contour $w_p w = w_a \lambda_d \lambda_s^{-1}$ as w_a slowly adjusts. In this fashion, the joint fast dynamics of w and w_p , combined with the slow dynamics of w_a , leads to a nonzero fixed point for all 3 values, despite the existence of a collapsed fixed point at the origin. Moreover, the larger the ratio $\lambda_d \lambda_s^{-1}$, which is determined by the data and augmentation, the larger the final values of both w and w_p will tend to be.

We can obtain further insight by noting that the submanifold $w = w_p$, in which the online and predictor weights are tied, constitutes an invariant submanifold of the dynamics in Eqns. 98 to 100; if $w = w_p$ at any instant of time, then this condition holds for all future time. Therefore we can both analyze and visualize the dynamics on this two dimensional invariant submanifold, with coordinates $w = w_p$ and w_a (Fig. 10 (middle)). This analysis clearly shows an unstable collapsed solution at the origin, with $w = w_a = 0$, and a stable non-collapsed solution at $w = w_a = \lambda_d \lambda_s^{-1}$.

We note again, that the generic existence of these non-collapsed solutions in Fig. 10 depends critically on the presence of a predictor with adjustable weights w_p . Removing the predictor corresponds to forcing $w_p = 1$, and non-collapsed solutions cannot exist unless $\lambda_d = \lambda_s$, as demonstrated in Fig. 10 (right). Thus, remarkably, in BYOL in this simple setting, the introduction of a predictor network plays a crucial role, even though it neither adds to the expressive capacity of the online network, nor improves its ability to match the target network. Instead, it plays a crucial role by dramatically modifying the learning dynamics (compare e.g. Fig 10 middle and right panels), thereby enabling convergence to noncollapsed solutions through a dynamical mechanism whereby the online and predictor network cooperatively amplify each others' weights to escape collapsed solutions (Fig. 10 (left)).

Overall, this analysis of BYOL learning dynamics provides considerable insight into the dynamical mechanisms enabling BYOL to avoid collapsed solutions, without negative pairs to force apart representations, in what is likely to be the simplest nontrivial setting. Further analysis on this model, in direct analogy to the analysis performed on the equivalent 3 dynamical system (derived under different assumptions) studied in the main paper, can yield similar insights into the dynamics of BYOL and SimSiam under various conditions on learning rates.