
An eigenspace view reveals how predictor networks and stop-grads provide implicit variance regularization

Manu Srinath Halvagal^{1,2,*}

Axel Laborieux^{1,*}

Friedemann Zenke^{1,2}

{firstname.lastname}@fmi.ch

¹ Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland

² Faculty of Natural Sciences, University of Basel, Basel, Switzerland

Abstract

Self-supervised learning (SSL) learns useful representations from unlabelled data by training networks to be invariant to pairs of augmented versions of the same input. Non-contrastive methods avoid collapse either by directly regularizing the covariance matrix of network outputs or through asymmetric loss architectures, two seemingly unrelated approaches. Here, by building on DirectPred [1], we lay out a theoretical framework that reconciles these two views. We derive analytical expressions for the representational learning dynamics in linear networks. By expressing them in the eigenspace of the embedding covariance matrix, where the solutions decouple, we reveal the mechanism and conditions that provide implicit variance regularization. These insights allow us to formulate a new isotropic loss function that equalizes eigenvalue contribution and renders learning more robust. Finally, we show empirically that our findings translate in nonlinear networks trained on CIFAR-10 and STL-10.

1 Introduction and background

SSL in Siamese architectures allows learning good representations without any labeled data [2–7]. These methods learn representations that are predictive across randomized transformations, i.e. the network’s objective is to “pull” together its outputs for two differently augmented versions of the same input. In order to avoid “representational collapse”, the trivial solution whereby network output becomes constant, SSL methods use either a contrastive objective to “push” apart representations of unrelated images [4, 5], explicit regularization of the representational covariance matrix [6, 7], or Siamese networks with an asymmetric loss [2, 3]. The latter strategy is used in the BYOL and SimSiam models. It is intriguing because it relies on a simple manipulation of the architecture and the flow of gradients in the network without any additional objective or regularization. Symmetry between the Siamese branches is broken by passing one of the representations through an additional predictor network and blocking gradients from flowing through the other “target” branch. The learning dynamics induced by such asymmetric loss Siamese architectures are surprisingly intricate [1, 8, 9] but not fully understood. A critical insight was provided by Tian et al. [1] for the case of linear feedforward and predictor networks. They showed that throughout training with the SimSiam loss, the eigenspace of the predictor weights aligns with that of the correlation matrix of the representations. Motivated by this insight, they proposed DirectPred, a simple strategy to directly set the predictor as a function of the estimated correlation matrix instead of learning it via gradient

* Equal contribution.

descent. Crucially, DirectPred and the further developed DirectCopy approach [8] remain effective at scale and outperform equivalent architectures with a linear learned predictor network.

In this work, we consider the linear network setting used previously in the theoretical analysis of non-contrastive SSL in [1, 8], and provide a straightforward analysis of how the learning dynamics in DirectPred and DirectCopy prevent representational collapse. Specifically, we show that these methods optimize an implicit loss function in the eigenspace of the representation correlation matrix that ensures the learning dynamics always converge to non-collapsed solutions. Furthermore, our analysis details out a clear connection to the covariance-regularization-based non-contrastive methods such as VICReg [7]. DirectPred and DirectCopy similarly regularize the variance of the representations, the only difference being that they do this in eigenspace instead of directly in the representation space as in VICReg. Finally, our analysis reveals that DirectPred may suffer from slow learning rates for components with associated small eigenvalues which is also a possible reason for why these approaches require an exponentially moving average (EMA) target network. Based on our analysis we propose an alternative isotropic loss function that mitigates this problem and is robust to removing the EMA.

2 Analysis of representational learning dynamics in Siamese networks

To gain analytical insight into how predictor networks and asymmetric loss configurations jointly prevent collapse, we consider a Siamese neural network $z = f(x; \theta)$ with $z \in \mathbb{R}^M$, the input $x \in \mathbb{R}^N$, and the parameters θ . We further assume a linear predictor network $W_P \in \mathbb{R}^{M \times M}$ and use the same parameters for the online and target branches as in SimSiam [3]. We consider pairs of embeddings $z^{(1/2)}$ corresponding to two inputs $x^{(1/2)}$ related through augmentation. Further we define $z = z^{(1)}$ and the correlation matrix of network outputs $\mathbb{E}_x [zz^\top]$ where the expectation is taken over the data distribution. Since the correlation matrix is a real symmetric matrix, we can diagonalize it as $\mathbb{E}_x [zz^\top] = UDU^\top$ where U is the orthogonal matrix whose columns are the eigenvectors of $\mathbb{E}_x [zz^\top]$ and D the diagonal matrix containing the eigenvalues λ_m with $m \in [1, M]$. Finally, instead of gradient-based optimization, we directly set the predictor network as a simple function of this correlation matrix as suggested in DirectPred [1]. Specifically,

$$W_P = f_\alpha (\mathbb{E}_x [zz^\top]) = UD^\alpha U^\top, \quad (1)$$

where α is a positive constant. We now consider the SimSiam loss $\mathcal{L}_{\text{SimSiam}}$ for a single pair of embeddings $z^{(1)}$ and $z^{(2)}$ and re-write it in the eigenbasis of W_P :

$$\begin{aligned} \mathcal{L}_{\text{SimSiam}} &= \frac{1}{2} \|W_P z^{(1)} - \text{SG}(z^{(2)})\|^2 \\ &= \frac{1}{2} \|UD^\alpha U^\top z^{(1)} - \text{SG}(UU^\top z^{(2)})\|^2 \\ &= \frac{1}{2} \|D^\alpha \hat{z}^{(1)} - \text{SG}(\hat{z}^{(2)})\|^2 \\ &= \frac{1}{2} \sum_m^M |\lambda_m^\alpha \hat{z}_m^{(1)} - \text{SG}(\hat{z}_m^{(2)})|^2 \end{aligned} \quad (2)$$

where we used the fact that U is orthogonal and therefore does not change the Euclidean norm. This transformation decouples the SimSiam loss into a sum of per-component losses over the eigenmodes.

We want to gain a theoretical understanding of representational changes during learning in the eigenbasis of W_P . To that end, we consider the continuous time gradient descent dynamics of $\mathcal{L}_{\text{SimSiam}}$ in the transformed space:

$$\dot{\hat{z}} = \nabla_\theta \dot{\hat{\theta}} = -\eta \hat{\Theta}_t(x, \mathcal{X}) \nabla_{\hat{z}} \mathcal{L} = -\eta \hat{\Theta}_t(x, \mathcal{X}) \left(D_t^\alpha \hat{z}_t^{(1)} - \hat{z}_t^{(2)} \right) D_t^\alpha, \quad (3)$$

where we introduced the the empirical neural tangent kernel (NTK) [10] $\hat{\Theta}_t(\mathcal{X}, \mathcal{X}) = \nabla_\theta \hat{z} \nabla_\theta \hat{z}^\top$ with the transformed representations $\hat{Z} = \hat{z}_t(\mathcal{X}) = U^\top z_t(\mathcal{X})$ of the whole dataset \mathcal{X} .

While $\hat{\Theta}_t$ changes over time and is generally intractable in finite-width networks it is always positive semidefinite. This property guarantees that the cosine angle between the representational training dynamics $\dot{\hat{Z}} \propto -\hat{\Theta}_t \nabla_{\hat{z}} \mathcal{L}$ and the hypothetical dynamics that result from optimizing the representations

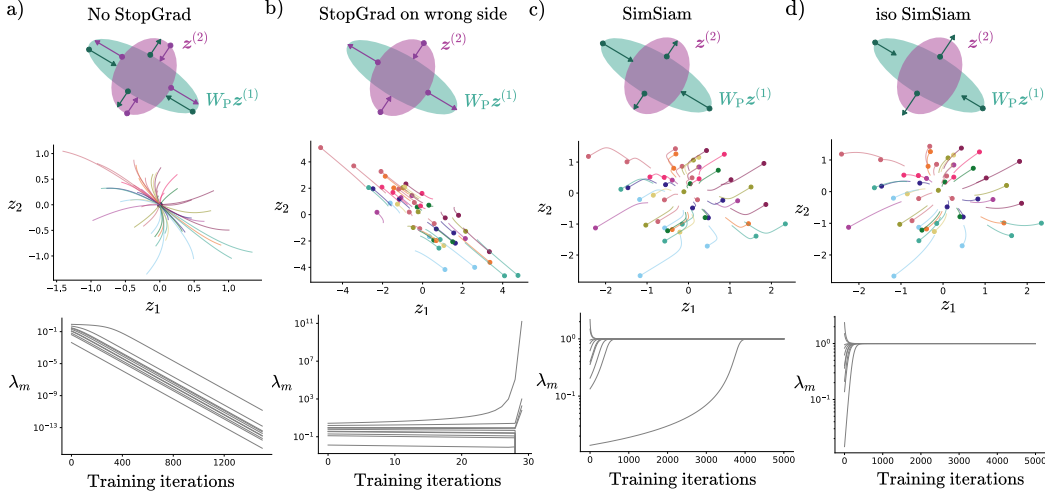


Figure 1: Top and middle rows show the neural updates in different settings, in dimension $M = 2$ for visualization. Bottom row shows the evolution of the eigenvalues of W_P upon training in the settings corresponding to the top row, but in dimensions $N = 15$ and $M = 10$. **a)** Omitting the stop grad leads to representational collapse. **b)** Applying the Stop grad on the wrong side also leads to collapse with potentially diverging eigenmodes. **c)** Optimizing the SimSiam loss leads to isotropic representations. **d)** Optimizing the isotropic loss has the same effect, but uniform for all eigenvalues.

directly $\dot{\hat{Z}} \propto -\nabla_{\hat{Z}} \mathcal{L}$ is always positive:

$$\left\langle -\nabla_{\hat{Z}} \mathcal{L}, \dot{\hat{Z}} \right\rangle = \eta \left\langle \nabla_{\hat{Z}} \mathcal{L}, \hat{\Theta}_t \nabla_{\hat{Z}} \mathcal{L} \right\rangle > 0. \quad (4)$$

Further, in the special case of linear networks with i.i.d Gaussian inputs, the NTK reduces to the identity resulting in perfect alignment (see Appendix A.1 for a proof). Despite the lack of perfect alignment in nonlinear networks, it served as our rationale to study representational dynamics in the idealized setting $\dot{\hat{Z}} \propto -\nabla_{\hat{Z}} \mathcal{L}$. In the idealized setting, representational updates decreasing Eq. (2) decouple $\dot{\hat{z}}_m^{(1)} = -\eta \frac{\partial \mathcal{L}}{\partial \hat{z}_m^{(1)}} = \eta \lambda_m^\alpha \left(\hat{z}_m^{(2)} - \lambda_m^\alpha \hat{z}_m^{(1)} \right)$ with learning rate η and in consequence each eigenmode evolves independently in expectation value:

$$\dot{\hat{z}}_m^{(1)} = \dot{\hat{z}}_m = \eta \lambda_m^\alpha \left(\bar{\hat{z}}_m^{(2)} - \lambda_m^\alpha \bar{\hat{z}}_m^{(1)} \right) = \eta \lambda_m^\alpha (1 - \lambda_m^\alpha) \hat{z}_m. \quad (5)$$

where we have omitted the time subscript t in \hat{z}_m and λ_m , and $\bar{(\cdot)}$ is the expectation taken over many random augmentations so that $\bar{\hat{z}}_m^{(1)} = \bar{\hat{z}}_m^{(2)} = \hat{z}_m$. From this expression, we appreciate that $\dot{\hat{z}}_m$ has the same sign as \hat{z}_m whenever $\lambda_m < 1$ and the opposite sign whenever $\lambda_m > 1$, which would subsequently drive the eigenvalue towards one, thereby preventing collapse of the representation.

In contrast, $\dot{\hat{z}}_m$ always has the opposite sign as \hat{z}_m for a loss without the StopGrad function:

$$\mathcal{L}_{\text{noSG}} := \frac{1}{2} \|W_P \mathbf{z}^{(1)} - \mathbf{z}^{(2)}\|^2 \Rightarrow \dot{\hat{z}}_m = -\eta (1 - \lambda_m^\alpha)^2 \hat{z}_m, \quad (6)$$

which is notorious for causing collapse [3].

Finally, the decoupled loss formulation also predicts the representational behavior when using StopGradient operation on the “wrong” side, i.e., on the online branch instead of the target branch. This manipulation is possible for DirectPred because the predictor needs no optimizing through gradient descent. In this case, the representational dynamics are governed by the following equations:

$$\mathcal{L}_{\text{ws}} := \frac{1}{2} \|\text{SG}(W_P \mathbf{z}^{(1)}) - \mathbf{z}^{(2)}\|^2 \Rightarrow \dot{\hat{z}}_m = \eta (\lambda_m^\alpha - 1) \hat{z}_m, \quad (7)$$

which leads to collapse for components with $\lambda_m < 1$ and divergence when $\lambda_m > 1$ (Fig. 1b)).

To verify these findings numerically, we simulated a small linear Siamese neural network with input dimension $N = 6$ and representation dimension $M = 2$. We fed the network with isotropic

Gaussian inputs, and pairs of augmentations were generated using isotropic Gaussian perturbations of standard deviation $\sigma = 0.1$. We then trained the linear encoder with the three different loss functions detailed above. Training the network with $\mathcal{L}_{\text{collapse}}$ from Eq. (6) resulted in collapse with exponentially decaying eigenvalues as predicted (Fig. 1a)). Optimizing \mathcal{L}_{ws} indeed resulted in diverging components with initial $\lambda_m > 1$ and collapse along other dimensions for which $\lambda_m < 1$ at initialization (Fig. 1b)) as predicted by our model. Finally, optimizing $\mathcal{L}_{\text{SimSiam}}$, the representations become increasingly isotropic with all the eigenvalues λ_m converging to one (Fig. 1c)), as predicted by Eq. (5). Importantly, these findings were qualitatively similar in nonlinear networks (see Fig. 2 in the Appendix).

In summary, our eigenspace analysis gives essential insights into the underlying mechanisms of representation learning in Siamese networks with predictor networks and sensible positioning of stop gradients in their loss formulation. First, it illustrates that representational dynamics implicitly drive the eigenvalues of the covariance matrix towards one thereby preventing complete collapse. Second, because collapse is prevented independently for all eigenvalues, the learning dynamics also effectively prevent dimensional collapse [9]. Together, these findings uncover a surprising link to covariance-regularization strategies, which enforce finite output variance and decorrelation [6, 7].

3 Isotropic loss functions speed up and stabilize self-supervised learning

We noted that the current eigenvalues act as effective learning rate modifiers in Eq. (5), meaning that large eigenvalues converge faster than smaller ones, reminiscent of previous theoretical work [11]. We speculated that this effect could lead to slow convergence for small eigenvalue components. However, using our framework it is straightforward to craft alternative ‘‘isotropic loss functions’’ that equalize relaxation dynamics for all eigenmodes. One such loss function is the following:

$$\mathcal{L}_{\text{iso}} = \frac{1}{2} \|\mathbf{z}^{(1)} - \text{SG}(\mathbf{z}^{(2)} + \mathbf{z}^{(1)} - W_P \mathbf{z}^{(1)})\|^2 \quad (8)$$

with associated learning dynamics $\hat{z}_m = \eta(1 - \lambda_m^\alpha) \hat{z}_m$. When optimizing \mathcal{L}_{iso} it resulted in similar convergence properties as the SimSiam loss, but all eigenmodes converged at more similar rates (Fig. 1d)) in agreement with our analysis.

Isotropic losses allow dispensing with moving average parameters in the target network. Previous work by Tian et al. [1] underscored the importance of boosting small eigenvalues and suggested that evolving the weights in the target network using EMA serves as an automatic curriculum for small eigenvalues. We speculated that the EMA may no longer be necessary when training with isotropic losses. To test this idea we trained a deep Siamese network on CIFAR-10 and STL-10 [12] using either the standard SimSiam loss or the isotropic loss introduced in Eq. (8). For these simulations, we used a VGG11 [13] backbone with a one-hidden-layer projection MLP, and SimCLR augmentations [4] using PyTorch Lightning [14, 15].

We first checked that training with IsoLoss \mathcal{L}_{iso} did not impair training accuracy in settings using an EMA target network. In fact, we found that IsoLoss resulted in slightly better accuracy on STL-10 compared to the regular SimSiam loss. Next we trained the same setup without EMA target network and found that IsoLoss outperformed the SimSiam loss by a large margin.

Table 1: Linear readout validation accuracies \pm stddev over four random seeds.

| | CIFAR-10 (800 epochs) | | STL-10 (100 epochs) | |
|--------------|-----------------------|----------------|---------------------|----------------|
| | EMA | no EMA | EMA | no EMA |
| SimSiam loss | 74.0 \pm 0.6 | 59.9 \pm 2.1 | 63.8 \pm 1.2 | 33.0 \pm 1.9 |
| IsoLoss | 72.5 \pm 0.4 | 73.5 \pm 0.5 | 67.1 \pm 0.4 | 65.7 \pm 0.8 |

4 Discussion

Building on DirectPred [1], we presented an eigenspace formulation that illustrates how representational collapse is avoided in Siamese SSL approaches with asymmetric loss configurations. Specifically, we showed that these architectures induce an implicit loss that regularizes the variance

and prevents dimensional collapse. Further, we showed that this formulation allows easily crafting new types of losses, such as IsoLoss that yield faster and more stable learning dynamics, thereby allowing to dispense with EMA target networks. Finally, the eigenspace view laid out in this article builds a conceptual bridge to other SSL approaches that explicitly regularize representational variance through additional loss terms [6, 7] and lays the foundation for better understanding modular architecture choices in terms of their equivalent objective functions.

References

- [1] Yuandong Tian, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning dynamics without contrastive pairs. In *International Conference on Machine Learning*, pages 10268–10278. PMLR, 2021.
- [2] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [3] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020.
- [6] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021.
- [7] Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [8] Xiang Wang, Xinlei Chen, Simon S Du, and Yuandong Tian. Towards demystifying representation learning with non-contrastive self-supervision. *arXiv preprint arXiv:2110.04947*, 2021.
- [9] Xiao Wang, Haoqi Fan, Yuandong Tian, Daisuke Kihara, and Xinlei Chen. On the importance of asymmetry for siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16570–16579, 2022.
- [10] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [11] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv:1312.6120 [cond-mat, q-bio, stat]*, December 2013. URL <http://arxiv.org/abs/1312.6120>.
- [12] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [13] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [14] William Falcon et al. Pytorch lightning. *GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>*, 3, 2019.
- [15] William Falcon and Kyunghyun Cho. A framework for contrastive self-supervised learning and designing a new approach. *arXiv preprint arXiv:2009.00104*, 2020.

A Appendix

A.1 The NTK reduces to the identity for a linear network with high-dimensional i.i.d Gaussian inputs

For a linear network with feedforward weights W , we first note that:

$$\begin{aligned}\hat{\mathbf{z}} &= U^\top f(\mathbf{x}) = U^\top W \mathbf{x} \\ \implies \nabla_\theta \hat{\mathbf{z}} &= \nabla_W \hat{\mathbf{z}} = \nabla_W (U^\top W \mathbf{x}) = \mathbf{x}^\top \otimes U^\top,\end{aligned}\tag{9}$$

where \otimes is the Kronecker product, resulting from the fact that every input component appears in the update once for each output component. Now, we begin our analysis with the gradient flow equation from Eq. (3):

$$\dot{\hat{\mathbf{z}}} = \nabla_\theta \hat{\mathbf{z}} \dot{\boldsymbol{\theta}} = -\eta \hat{\Theta}_t(\mathbf{x}, \mathcal{X}) \nabla_{\hat{\mathbf{z}}} \mathcal{L}\tag{10}$$

where $\hat{\Theta}_t(\mathbf{x}, \mathcal{X}) = \nabla_\theta \hat{\mathbf{z}} \nabla_\theta \hat{\mathbf{z}}^\top$ is the $(M \times M|\mathcal{D}|)$ block row of the full $(M|\mathcal{D}| \times M|\mathcal{D}|)$ empirical NTK $\hat{\Theta}_t(\mathcal{X}, \mathcal{X})$ that corresponds to a single data sample \mathbf{x} . The diagonal blocks in $\hat{\Theta}_t(\mathcal{X}, \mathcal{X})$ correspond to single samples and the off-diagonal blocks are cross-terms between samples from the full dataset. We can develop a generic expression for each block $\hat{\Theta}_t(\mathbf{x}_i, \mathbf{x}_j)$ corresponding to the interactions between samples i and j as:

$$\begin{aligned}\hat{\Theta}_t(\mathbf{x}_i, \mathbf{x}_j) &= \nabla_W \hat{\mathbf{z}}_i \nabla_W \hat{\mathbf{z}}_j^\top \\ &= (\mathbf{x}_i^\top \otimes U^\top) (\mathbf{x}_j^\top \otimes U^\top)^\top \\ &= (\mathbf{x}_i^\top \otimes U^\top) (\mathbf{x}_j \otimes U) \\ &= (\mathbf{x}_i^\top \mathbf{x}_j) \otimes (U^\top U) \\ &= (\mathbf{x}_i^\top \mathbf{x}_j) \otimes I_M \\ &= (\mathbf{x}_i^\top \mathbf{x}_j) I_M.\end{aligned}\tag{11}$$

where we have used the fact that $(A \otimes B)^\top = A^\top \otimes B^\top$ and $(A \otimes B)(C \otimes D) = AC \otimes BD$. Here, I_M is the identity matrix of size M . We note here that Eq. (11) constitutes a general result that the NTK for a linear network is invariant under orthogonal transformations of the network output.

We now calculate the average interaction between samples drawn from an i.i.d standard Gaussian distribution. In this case, for high-dimensional inputs \mathbf{x} we have $\mathbf{x}_i^\top \mathbf{x}_j \approx \delta_{ij}$. Finally, we see that:

$$\begin{aligned}\dot{\hat{\mathbf{z}}}_i &= -\eta \hat{\Theta}_t(\mathbf{x}_i, \mathcal{X}) \nabla_{\hat{\mathbf{z}}} \mathcal{L} \\ &= -\eta \hat{\Theta}_t(\mathbf{x}_i, \mathbf{x}_i) \nabla_{\hat{\mathbf{z}}_i} \mathcal{L} - \eta \sum_{j \neq i} \hat{\Theta}_t(\mathbf{x}_i, \mathbf{x}_j) \nabla_{\hat{\mathbf{z}}_j} \mathcal{L} \\ &= -\eta (\mathbf{x}_i^\top \mathbf{x}_i) \nabla_{\hat{\mathbf{z}}_i} \mathcal{L} - \eta \sum_{j \neq i} (\mathbf{x}_i^\top \mathbf{x}_j) \nabla_{\hat{\mathbf{z}}_j} \mathcal{L} \\ &= -\eta \nabla_{\hat{\mathbf{z}}_i} \mathcal{L}\end{aligned}\tag{12}$$

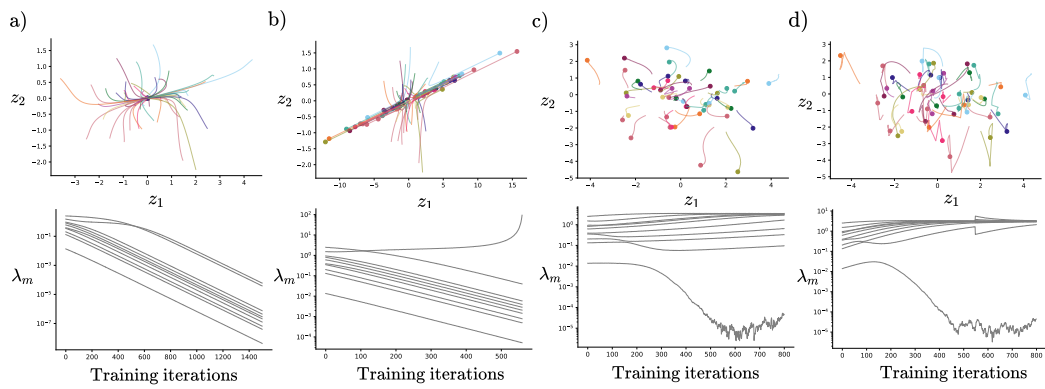


Figure 2: Same as Fig. 1, but for a nonlinear network using rectified linear units.