

A THEORY OF UNIMODAL BIAS IN MULTIMODAL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Using multiple input streams simultaneously in training multimodal neural networks is intuitively advantageous, but practically challenging. A key challenge is unimodal bias, where a network overly relies on one modality and ignores others during joint training. While unimodal bias is well-documented empirically, our theoretical understanding of how architecture and data statistics influence this bias remains incomplete. Here we develop a theory of unimodal bias with deep multimodal linear networks. We calculate the duration of the unimodal phase in learning as a function of the depth at which modalities are fused within the network, dataset statistics, and initialization. We find that the deeper the layer at which fusion occurs, the longer the unimodal phase. In addition, our theory reveals the modality learned first is not necessarily the modality that contributes more to the output. Our results, derived for multimodal linear networks, extend to ReLU networks in certain settings. Taken together, this work illuminates pathologies of multimodal learning under joint training, showing that late and intermediate fusion architectures can give rise to long unimodal phases and even prioritize learning a less helpful modality.

1 INTRODUCTION

The success of multimodal deep learning hinges on effectively utilizing multiple modalities (Baltrušaitis et al., 2019; Liang et al., 2022). However, some multimodal networks overly rely on a faster-to-learn or easier-to-learn modality and ignore the others during joint training (Goyal et al., 2017; Cadene et al., 2019; Wang et al., 2020; Gat et al., 2021; Peng et al., 2022). For example, Visual Question Answering models should provide a correct answer by both “listening” to the question and “looking” at the image (Agrawal et al., 2016), whereas they tend to overly rely on the language modality and ignore the visual modality (Goyal et al., 2017; Agrawal et al., 2018; Hessel & Lee, 2020). This phenomenon has been observed in a variety of settings, and has several names: unimodal bias (Cadene et al., 2019), greedy learning (Wu et al., 2022), modality competition (Huang et al., 2022), modality laziness (Du et al., 2023), and modality underutilization (Makino et al., 2023).

In this paper we adopt the term unimodal bias to refer to the phenomenon in which a multimodal network learns from different input modalities at different times during joint training. The extent to which multimodal networks exhibit unimodal bias depends on both the dataset and the multimodal network. Goyal et al. (2017); Agrawal et al. (2018); Hudson & Manning (2019) tried to alleviate the bias by building more balanced multimodal datasets. Empirical work has shown that unimodal bias emerges in jointly trained late fusion networks (Wang et al., 2020; Huang et al., 2022) and intermediate fusion networks (Wu et al., 2022), while early fusion networks may encourage usage of all input modalities (Gadzicki et al., 2020; Barnum et al., 2020).

Despite empirical evidence, there is scarce theoretical understanding of how unimodal bias arises and how it is affected by the network configuration, dataset statistics, and initialization. To unravel this pathological behavior, we study deep multimodal linear networks with three common fusion schemes: early, intermediate, and late (Ramachandram & Taylor, 2017). We find that unimodal bias is absent in early fusion linear networks but present in intermediate and late fusion linear networks. Intermediate and late fusion linear networks learn one modality first and the other after a delay, yielding a phase in which the multimodal network implements a unimodal function. This difference in time is consistent with the empirical observation that multimodal networks learn different

modalities at different speeds (Wang et al., 2020; Wu et al., 2022). We compute the duration of the unimodal phase in terms of parameters of the network and the dataset. We find that a deeper fusion layer within the multimodal network, stronger correlations between input modalities, and greater disparities in input-output correlations for each modality all prolong the unimodal phase. We also find that multimodal networks have a superficial preference for which modality to learn first: they prioritize the faster-to-learn modality, which is not necessarily the modality that contributes more to the output, and we derive conditions under which a network will exhibit this superficial preference. Our results apply to ReLU networks in certain settings, providing insights for examining, diagnosing, and curing unimodal bias in a broader range of realistic cases.

Our contributions are the following: (i) We provide a theoretical explanation for the presence of unimodal bias in late and intermediate fusion linear networks and the absence of unimodal bias in early fusion linear networks. (ii) We calculate the duration of the unimodal phase in multimodal learning with late and intermediate fusion linear networks, as a function of the network configuration, correlation matrices of the dataset, and initialization scale. (iii) We show that under specific conditions on dataset statistics, late and intermediate fusion linear networks exhibit a superficial modality preference, prioritizing learning a modality that makes less contribution to the output. (iv) We validate our findings with numerical simulations of deep linear networks and certain two-layer ReLU networks.

1.1 RELATED WORK

Several attempts have been made to explain the unimodal bias behavior. Various metrics have been proposed to inspect the internal mechanics of multimodal learning and multimodal models (Wang et al., 2020; Wu et al., 2022; Kleinman et al., 2023). Huang et al. (2021; 2022) provides a theoretical explanation for why multimodal learning has the capacity to outperform unimodal learning but can fail to deliver. Makino et al. (2023) propose incidental correlation to diagnose and explain modality underutilization in the small data regime. Our work differs by seeking an analytical relationship between unimodal bias, network configuration, and dataset statistics.

Our work leverages a rich line of theoretical literature on deep linear neural networks. Exact solutions and reductions of the gradient descent dynamics have been derived in (Fukumizu, 1998; Saxe et al., 2014; 2019; Arora et al., 2018; Advani et al., 2020; Atanasov et al., 2022; Shi et al., 2022). Balancing properties in linear networks were discovered and proved in (Ji & Telgarsky, 2019; Du et al., 2018). However, as multimodal linear networks are not fully connected and multimodal datasets generally do not have white input covariance, previous solutions no longer apply, and we had to develop new analytic tools.


2 PROBLEM SETUP

2.1 MULTIMODAL DATA

Let $\mathbf{x} \in \mathbb{R}^D$ represent an arbitrary multimodal input and $y \in \mathbb{R}$ be its scalar target output. We are given a dataset $\{\mathbf{x}^\mu, y^\mu\}_{\mu=1}^P$ consisting of P samples. For simplicity, we assume there are two modalities A and B with full input $\mathbf{x} = [\mathbf{x}_A, \mathbf{x}_B]^\top$. Since we study multimodal linear networks, the learning dynamics only depend on correlation matrices of the dataset (Fukumizu, 1998; Saxe et al., 2014). We notate the input correlation matrix as Σ and input-output correlation matrix as $\Sigma_{y\mathbf{x}}$ defined as

$$\Sigma = \begin{bmatrix} \Sigma_A & \Sigma_{AB} \\ \Sigma_{BA} & \Sigma_B \end{bmatrix} = \begin{bmatrix} \langle \mathbf{x}_A \mathbf{x}_A^\top \rangle & \langle \mathbf{x}_A \mathbf{x}_B^\top \rangle \\ \langle \mathbf{x}_B \mathbf{x}_A^\top \rangle & \langle \mathbf{x}_B \mathbf{x}_B^\top \rangle \end{bmatrix}, \Sigma_{y\mathbf{x}} = [\Sigma_{y\mathbf{x}_A} \quad \Sigma_{y\mathbf{x}_B}] = [\langle y \mathbf{x}_A^\top \rangle \quad \langle y \mathbf{x}_B^\top \rangle] \quad (1)$$

where $\langle \cdot \rangle$ denotes the average over the dataset. We assume data points are centered to have zero mean $\langle \mathbf{x} \rangle = \mathbf{0}$ and covariance Σ has full rank, but make no further assumptions on the correlation matrices.

Figure 1: Schematic of a multimodal fusion network with total depth L and fusion layer at L_f .

2.2 MULTIMODAL FUSION LINEAR NETWORK

We study a deep linear network with total depth L and fusion layer at L_f defined as

$$\hat{y}(\mathbf{x}; \mathbf{W}) = \prod_{j=L_f+1}^L \mathbf{W}^j \left(\prod_{i=1}^{L_f} \mathbf{W}_A^i \mathbf{x}_A + \prod_{i=1}^{L_f} \mathbf{W}_B^i \mathbf{x}_B \right) \equiv \mathbf{W}_A^{\text{tot}} \mathbf{x}_A + \mathbf{W}_B^{\text{tot}} \mathbf{x}_B \equiv \mathbf{W}^{\text{tot}} \mathbf{x}. \quad (2)$$

The overall network input-output map is denoted as \mathbf{W}^{tot} and the map for each modality is denoted as $\mathbf{W}_A^{\text{tot}}, \mathbf{W}_B^{\text{tot}}$. We use \mathbf{W} to denote all weight parameters collectively. We assume the number of neurons in both pre-fusion layer branches is of the same order. A schematic of this network is given in Fig. 1.

Our network definition incorporates bimodal deep linear networks of three common fusion schemes. Categorized by the multimodal deep learning community (Ramachandram & Taylor, 2017), the case where $L_f = 1$ is early or data-level fusion; $1 < L_f < L$ is intermediate fusion; and $L_f = L$ is late or decision-level fusion.

2.3 GRADIENT DESCENT DYNAMICS

The network is trained by optimizing the mean square error loss $\mathcal{L} = \frac{1}{2P} \sum_{\mu=1}^P (\hat{y}^\mu - y^\mu)^2$ with full batch gradient descent. In the limit of small learning rate, the gradient descent dynamics are well approximated by the continuous time differential equations given below; see Appendix A. For pre-fusion layers $1 \leq l \leq L_f$,

$$\tau \dot{\mathbf{W}}_A^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_A^i \right)^\top (\boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^{\text{tot}} \boldsymbol{\Sigma}_A - \mathbf{W}_B^{\text{tot}} \boldsymbol{\Sigma}_{BA}) \left(\prod_{i=1}^{l-1} \mathbf{W}_A^i \right)^\top, \quad (3a)$$

$$\tau \dot{\mathbf{W}}_B^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_B^i \right)^\top (\boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_A^{\text{tot}} \boldsymbol{\Sigma}_{AB} - \mathbf{W}_B^{\text{tot}} \boldsymbol{\Sigma}_B) \left(\prod_{i=1}^{l-1} \mathbf{W}_B^i \right)^\top. \quad (3b)$$

For post-fusion layers $L_f + 1 \leq l \leq L$,

$$\begin{aligned} \tau \dot{\mathbf{W}}^l = & \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (\boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^{\text{tot}} \boldsymbol{\Sigma}_A - \mathbf{W}_B^{\text{tot}} \boldsymbol{\Sigma}_{BA}) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_A^i \right)^\top \\ & + \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (\boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_A^{\text{tot}} \boldsymbol{\Sigma}_{AB} - \mathbf{W}_B^{\text{tot}} \boldsymbol{\Sigma}_B) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_B^i \right)^\top, \end{aligned} \quad (4)$$

where the time constant τ is the inverse of learning rate $\eta^{-1} = \tau$. We abuse the notation $\prod_j \mathbf{W}^j$ to represent the ordered product of matrices with the largest index on the left and smallest on the right. The network is initialized with small random weights.

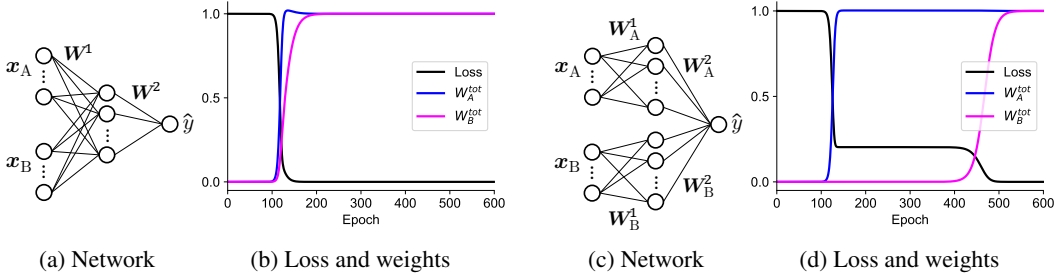


Figure 2: Schematic and training trajectories of two-layer early and late fusion linear networks. (a,c) Schematic of a two-layer early fusion (a) and late fusion (c) linear network. (b,d) Loss and total weights trajectories in the two-layer early fusion (b) and late fusion (d) linear network. Details: Both networks are trained to learn the same task of fitting $y = x_A + x_B$. Inputs x_A and x_B are scalars with covariance matrix $\Sigma = \text{diag}(4, 1)$. The early fusion network has 100 hidden neurons. The late fusion network has 100 hidden neurons in both branches. Both networks are initialized with small random weights sampled independently from $\mathcal{N}(0, 10^{-9})$. The learning rate is 0.04.

3 TWO-LAYER MULTIMODAL LINEAR NETWORKS

We first study two-layer multimodal linear networks with $L = 2$. There are two possible fusion schemes for two-layer networks according to our setup, namely early fusion $L_f = 1$, as in Fig. 2a and late fusion $L_f = 2$, as in Fig. 2c. Unimodal bias is absent in early fusion linear networks but present in late fusion linear networks. We first give an overview of the qualitative behavior based on the loss landscape and then quantitatively characterize the unimodal bias.

3.1 TWO-LAYER EARLY FUSION LINEAR NETWORK

Early fusion networks learn from both modalities simultaneously as shown in Fig. 2b. There is no conspicuous unimodal phase over the course of training.

This behavior can be appreciated through an analysis of the fixed points of the gradient dynamics. There are two manifolds of fixed points in early fusion networks (Appendix B): one is an unstable fixed point at zero \mathcal{M}_0 and the other is a manifold of stable fixed points at \mathcal{M}_* ,

$$\mathcal{M}_0 = \{\mathbf{W} | \mathbf{W} = \mathbf{0}\}; \tag{5a}$$

$$\mathcal{M}_* = \{\mathbf{W} | \mathbf{W}^{\text{tot}} = \Sigma_{y\mathbf{x}} \Sigma^{-1}\}. \tag{5b}$$

The network starts from small initialization, which is close to the unstable fixed point \mathcal{M}_0 . When learning progresses, the network escapes from the unstable fixed point \mathcal{M}_0 and converges to the global pseudo-inverse solution at \mathcal{M}_* . As studied by Saxe et al. (2019); Jacot et al. (2021); Pesme & Flammarion (2023), linear networks trained from small initialization exhibit quasi-stage-like behavior: learning progresses slowly for most of the time and rapidly moves from one fixed point or saddle to the next with a brief sigmoidal transition. Because there is only one fixed point aside from the zero fixed point at initialization to transit to, all dimensions and thus all modalities are learned during this one transition. Since the transitional period is very brief compared to the total learning time, all modalities are learned at almost the same time in early fusion networks.

3.2 TWO-LAYER LATE FUSION LINEAR NETWORK

Late fusion networks learn from two modalities with two sigmoidal transitions at two different times, as shown in Fig. 2d. There can be a prolonged unimodal phase during joint training.

Late fusion linear networks have the same two manifolds of fixed points $\mathcal{M}_0, \mathcal{M}_*$ as early fusion networks. In addition, late fusion linear networks have two manifolds of saddles $\mathcal{M}_A, \mathcal{M}_B$ (Appendix C.1), corresponding to learning one modality but not the other,

$$\mathcal{M}_A = \{\mathbf{W} | \mathbf{W}_A^{\text{tot}} = \Sigma_{y x_A} \Sigma_A^{-1}, \mathbf{W}_B^{\text{tot}} = \mathbf{0}\}; \tag{6a}$$

$$\mathcal{M}_B = \{\mathbf{W} | \mathbf{W}_A^{\text{tot}} = \mathbf{0}, \mathbf{W}_B^{\text{tot}} = \Sigma_{y x_B} \Sigma_B^{-1}\}. \tag{6b}$$

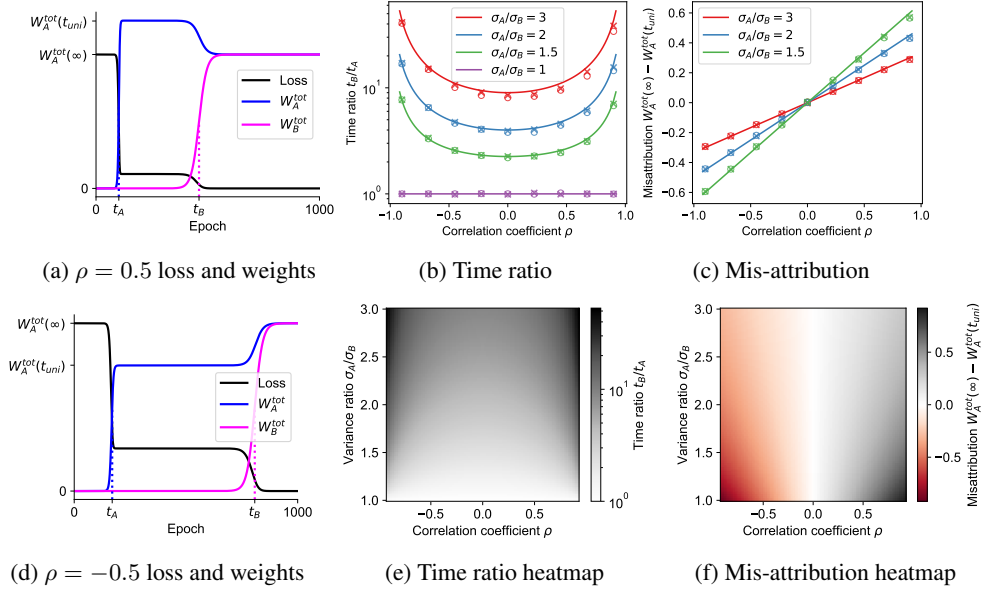


Figure 3: The duration of the unimodal phase and the amount of mis-attribution in two-layer late fusion linear networks. Two-layer late fusion linear networks of the same setting as Fig. 2c are trained to fit $y = x_A + x_B$ with different input covariance matrices. We denote the elements in the two-dimensional input covariance matrix as $\Sigma = [\sigma_A^2, \rho\sigma_A\sigma_B; \rho\sigma_A\sigma_B, \sigma_B^2]$. (a,d) Examples of loss and total weight trajectories in two-layer late fusion networks when modalities have positive correlation $\rho = 0.5$ (a) and negative correlation $\rho = -0.5$ (d). (b,c,e,f) Time ratio and amount of mis-attribution. Lines and heatmaps are theoretical predictions; circles are simulations of two-layer late fusion linear networks; crosses are simulations of two-layer late fusion ReLU networks. When assuming inputs x_A, x_B are scalars and the target output is $y = x_A + x_B$, the time ratio reduces to $1 + (\frac{\sigma_A^2}{\sigma_B^2} - 1)/(1 - \rho^2)$; the amount of mis-attribution reduces to $\rho \frac{\sigma_B}{\sigma_A}$.

The late fusion linear network therefore goes through two transitions because the network first arrives and lingers near a saddle in \mathcal{M}_A or \mathcal{M}_B and subsequently converges to the global pseudo-inverse solution in \mathcal{M}_* . Visiting the saddle gives rise to the plateau in the loss that separates the time when the two modalities are learned. We visualize the learning path and fixed points in the phase portrait in Fig. 6.

Which modality is learned first depends solely on the relative size of $\|\Sigma_{yx_A}\|$ and $\|\Sigma_{yx_B}\|$ in late fusion linear networks with small initialization¹. For definiteness, in all our simulations we choose $\|\Sigma_{yx_A}\| > \|\Sigma_{yx_B}\|$ which ensures that modality A is learned first; see Appendix C.3.1. We define t_A to be the time when the total weight of modality A reaches half of its associated plateau; and similarly for t_B as illustrated in Figs. 3a and 3d. At time t_A , W_A^{tot} escapes from the zero fixed point \mathcal{M}_0 and the network visits a saddle in \mathcal{M}_A where modality A fits the output as much as it can while modality B has not been learned. During the time lag from t_A to t_B , the loss lingers in a plateau and the network is unimodal. At time t_B , modality B escapes from the zero fixed point and the network converges to the global pseudo-inverse solution fixed point in \mathcal{M}_* . We quantify how long the unimodal phase is in Section 3.2.1. We take a closer look at the unimodal mode during the unimodal phase by highlighting the mis-attribution in unimodal phase in Section 3.2.2 and the superficial modality preference in Section 3.2.3.

3.2.1 DURATION OF THE UNIMODAL PHASE

We first quantify the duration of the unimodal phase. Because of the small random initialization, we can assume that the Frobenius norm of W_A^1 is approximately equal to the norm of W_B^1 at initialization, notated as u_0 . Through leading order approximation, we compute the times t_A and t_B in

¹We use $\|\cdot\|$ to notate the L2 norm of a vector or the Frobenius norm of a matrix in this paper.

Appendix C.3, which yields

$$t_A \approx \tau \|\Sigma_{yx_A}\|^{-1} \ln \frac{1}{u_0}, \quad t_B \approx t_A + \tau \frac{1 - \|\Sigma_{yx_A}\|^{-1} \|\Sigma_{yx_B}\|}{\|\Sigma_{yx_B} - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB}\|} \ln \frac{1}{u_0}. \quad (7)$$

To compare the unimodal phase duration across different settings, we focus on the time ratio,

$$\frac{t_B}{t_A} = 1 + \frac{\|\Sigma_{yx_A}\| - \|\Sigma_{yx_B}\|}{\|\Sigma_{yx_B} - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB}\|}. \quad (8)$$

We note that the time lag $t_B - t_A$ is proportional to $\|\Sigma_{yx_A}\| - \|\Sigma_{yx_B}\|$, which accords with the intuition that the $\|\Sigma_{yx_A}\|$ governs the speed at which modality A is learned and $\|\Sigma_{yx_B}\|$ governs the speed at which modality B is learned during time 0 to t_A . The denominator governs the speed at which modality B is learned during the unimodal phase from time t_A to t_B . Specifically, the network visits a saddle in \mathcal{M}_A during the unimodal phase, which affects the speed at which modality B is learned during the unimodal phase if cross correlation is not zero. Near the saddle, the speed of modality B is reduced by $\Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB}$; see Appendix C.3.

We validate Eq. (8) with numerical simulations in Fig. 3b. From Eq. (8) and Fig. 3b, we conclude that stronger correlations between input modalities and a greater disparity in input-output correlations for each modality make the time ratio larger, indicating a longer unimodal phase. In the extreme case of having maximum correlation, x_A and x_B have collinearity and one of them is redundant. Thus the denominator $\Sigma_{yx_B} - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB} = \mathbf{0}$ and the time ratio is ∞ , implying later becomes never — the network learns to fit the output only with modality A and modality B will never be learned as shown in Fig. 7.

We also simulate two-layer late fusion networks with ReLU nonlinearity added to the hidden layer, while holding the rest of our setting fixed. We empirically find that two-layer late fusion ReLU networks trained to learn a linear target map learn two modalities with two transitions and the time ratio plotted in Fig. 3b (crosses) closely aligns with the theoretical predictions derived for late fusion linear networks.

3.2.2 MIS-ATTRIBUTION IN THE UNIMODAL PHASE

We now quantify how much the multimodal network mis-attributes some of the output to modality A during the unimodal phase. When modalities are correlated, the local pseudo-inverse solution differs from the global pseudo-inverse solution ($[\Sigma_{yx_A} \Sigma_A^{-1}, \Sigma_{yx_B} \Sigma_B^{-1}] \neq \Sigma_{yx} \Sigma^{-1}$). During the unimodal phase, W_A^{tot} fits the output as much as it can and the network mis-attributes some of the output contributed by modality B to modality A by exploiting their correlations. Specifically, the weights of modality A overshoot if modalities have a positive correlation as in Fig. 3a and undershoot if negative as in Fig. 3d. This mis-attribution is then corrected when modality B catches up and the network eventually converges to the global pseudo-inverse solution. We demonstrate, using scalar input for clarity, that mis-attribution is more severe when modalities have higher correlation; see Figs. 3c and 3f.

When modalities are uncorrelated, late fusion networks do not mis-attribute during the unimodal phase. Weights for modality A converge to the global pseudo-inverse solution at time t_A and do not change thereafter, as in Fig. 2d, since the local pseudo-inverse solutions are the same as the global pseudo-inverse solution. In this case, late fusion networks behave the same as separately trained unimodal networks.

We empirically find that mis-attribution exists in two-layer late fusion ReLU networks as well and the amount of mis-attribution plotted in Fig. 3c (crosses) closely align with theoretical predictions for late fusion linear networks.

3.2.3 SUPERFICIAL MODALITY PREFERENCE

We now look into which modality is learned first. Late fusion networks have what we call “superficial modality preference” for which modality to learn first. They prioritize the modality that is faster to learn, which is not necessarily the modality that yields the larger decrease in loss. Put differently, late fusion networks first learn the modality that has a higher correlation with the output,




Figure 4: Demonstration of superficial modality preference. Two-layer late fusion linear networks with the same setting as Fig. 2 are trained with different datasets. (a,b) In both examples, modality A is learned first. The dotted black line marks the loss when the network visits \mathcal{M}_A . The dotted gray line marks the loss if the network had instead visited \mathcal{M}_B . (a) The target output is $y = x_A + 4x_B$ and the input covariance matrix is $\Sigma = \text{diag}(9, 1)$. The dotted gray line is below the dotted black, meaning modality B contributes more to the output. The prioritized modality is therefore not the modality that contributes more to the output. (b) The target output is $y = x_A + 3x_B$ and the input covariance matrix is $\Sigma = \text{diag}(16, 1)$. The dotted black line is below the dotted gray, meaning modality A contributes more to the output. The prioritized modality is the modality that contributes more to the output. (c) Boundaries of which modality is prioritized and which modality contributes more to the output in terms of dataset statistics. In region I and III, modality A is learned first. In region I and II, modality A contributes more to the output. Thus in region II and III (shaded red), prioritization and contribution disagree, resulting in superficial modality preference.

even though the other modality may make a larger contribution to the output. Under the following two conditions on dataset statistics,

$$\|\Sigma_{yx_A}\| > \|\Sigma_{yx_B}\|, \Sigma_{yx_A}\Sigma_A^{-1}\Sigma_{yx_A}^\top < \Sigma_{yx_B}\Sigma_B^{-1}\Sigma_{yx_B}^\top, \quad (9)$$

modality A is faster to learn but modality B contributes more to the output. We present an example where these conditions hold in Fig. 4a and where they do not in Fig. 4b.

If x_A, x_B are scalars and uncorrelated, the two inequality conditions in Eq. (9) reduce to

$$\frac{\sigma_B^2}{\sigma_A^2} < \frac{w_A^*}{w_B^*} < \frac{\sigma_B}{\sigma_A}, \quad (10)$$

where σ_A, σ_B are variances of x_A, x_B and we assume the target output is generated as $y = w_A^*x_A + w_B^*x_B$. We plot the two conditions in Fig. 4c. Region III satisfies the conditions we give in Eq. (10). Region II corresponds to Eq. (10) with flipped inequality signs, meaning the other superficial modality preference case where modality B is prioritized but modality A contributes more to the output. Hence, region II and III (shaded red) cover the dataset statistics where prioritization and contribution disagree and late fusion linear networks would prioritize learning the modality that contributes less to the output.

4 DEEP MULTIMODAL LINEAR NETWORKS

We now consider the more general case of deep multimodal linear networks and examine how the fusion layer depth L_f modulates the extent of unimodal bias.

4.1 DEEP EARLY FUSION LINEAR NETWORK

Similar to two-layer early fusion networks, deep early fusion networks learn from both modalities simultaneously, as shown in Fig. 5a (purple curve). Saxe et al. (2014; 2019); Advani et al. (2020) have shown that for deep linear networks, depth slows down learning but does not qualitatively change the dynamics compared to two-layer linear networks. The weights associated with all input modalities escape from the initial zero fixed point \mathcal{M}_0 and converge to the pseudo-inverse solution fixed point in \mathcal{M}_* in one transitional period.




Figure 5: Duration of the unimodal phase in deep multimodal linear networks. Four-layer linear networks with fusion layer at $L_f = 1, 2, 3, 4$ are trained to fit $y = \mathbf{x}_A + \mathbf{x}_B$. (a) An example of loss trajectories when input covariance matrix $\Sigma = \text{diag}(4, 1)$. (b,c,d) Lines are theory; circles are simulations of four-layer linear networks with different fusion layer depth. Variances of $\mathbf{x}_A, \mathbf{x}_B$ and their correlation coefficient are notated as σ_A, σ_B and ρ respectively. (b) Correlation coefficient sweep with $\sigma_A/\sigma_B = 2$ and initialization scale $u_0 = 0.1$. (c) Variance ratio sweep with $\rho = 0, u_0 = 0.1$. Note that $\sigma_A/\sigma_B = \sqrt{k}$ when $\rho = 0$. (d) Initialization scale sweep with $\sigma_A/\sigma_B = 2, \rho = 0$.

4.2 DEEP INTERMEDIATE AND LATE FUSION LINEAR NETWORK

Deep intermediate and late fusion linear networks learn the two modalities with two separate transitions, as shown in Fig. 5a; this is similar to what happens in two-layer late fusion networks. Due to the common terms in Eqs. (3) and (4) that govern the weight dynamics, the two manifolds of fixed points, $\mathcal{M}_0, \mathcal{M}_*$, and the two manifolds of saddles, $\mathcal{M}_A, \mathcal{M}_B$, still exist for any $2 \leq L_f \leq L, L \geq 2$, comprising intermediate and late fusion linear networks of any configuration.

In what follows, we stick to the convention that modality A is learned first. Deep intermediate and late fusion networks start from small initialization, which is close to the zero fixed point \mathcal{M}_0 . When modality A is learned in the first transitional period at time t_A , the network visits the saddle \mathcal{M}_A . After a unimodal phase, the network goes through the second transition at time t_B to reach the global pseudo-inverse solution fixed point \mathcal{M}_* . Because the network is in the same manifold \mathcal{M}_A during the unimodal phase, our results in Section 3.2.2 on the mis-attribution in the unimodal phase and Section 3.2.3 on superficial modality preference in two-layer late fusion networks directly carry over to deep intermediate and late fusion linear networks.

As shown in Fig. 5, the loss trajectories of networks with different $L_f \geq 2$ traverse the same plateau but they stay in the plateau for different durations. We thus quantify how the total depth L and fusion layer depth L_f affect the duration of the unimodal phase in Section 4.2.1.

4.2.1 DURATION OF THE UNIMODAL PHASE

We now calculate the duration of the unimodal phase in deep intermediate and late fusion linear networks, incorporating the new parameters L and L_f . The input-output correlation ratio is denoted as $k = \|\Sigma_{y\mathbf{x}_B}\|/\|\Sigma_{y\mathbf{x}_A}\| \in (0, 1)$. We derive the expression of the time ratio through leading order approximation in Appendix D. For $2 < L_f \leq L$, the time ratio is

$$\frac{t_B}{t_A} = 1 + \frac{(\|\Sigma_{y\mathbf{x}_A}\| - \|\Sigma_{y\mathbf{x}_B}\|)u_0^{L-L_f}}{(L_f - 2)\|\Sigma_{y\mathbf{x}_A}\Sigma_A^{-1}\|^{1-\frac{L_f}{L}}\|\Sigma_{y\mathbf{x}_B} - \Sigma_{y\mathbf{x}_A}\Sigma_A^{-1}\Sigma_{AB}\|} I(L, L_f)^{-1}, \quad (11)$$

where

$$I(L, L_f) = \int_1^\infty x^{1-L} \left[1 + (k + (1-k)x^{L_f-2})^{\frac{2}{2-L_f}} \right]^{\frac{L_f-L}{2}} dx. \quad (12)$$

For $L_f = 2$, the expression is slightly different; see Appendix D.3. As shown in Fig. 5, the theoretical prediction captures the trend that a deeper fusion layer L_f , a larger input-output correlation ratio k , and stronger correlations Σ_{AB} between input modalities all prolong the duration of the unimodal phase. The qualitative influence of dataset statistics on the time ratio in intermediate and late fusion networks is consistent with what we have seen in two-layer late fusion networks.

We now look into the influence of the fusion layer depth. By setting $L_f = L$ in Eq. (11), we find that the time ratio in deep late fusion networks reduces to the same expression as the two-layer late

fusion case in Eq. (8), which only involves dataset statistics but not the depth of the network or the initialization, since depth slows down the learning of both modalities by the same factor. In intermediate fusion linear networks, the time ratio is smaller than that in late fusion networks, with a smaller ratio with a shallower fusion layer. In intermediate fusion linear networks, learning one modality changes the weights in its associated pre-fusion layers and the shared post-fusion layers. At time t_A , the pre-fusion layer weights of modality A and the shared post-fusion layer weights have escaped from the zero fixed point and grown in scale while the pre-fusion layer weights of modality B have not. During the unimodal phase, the shared post-fusion layer weights and the correlation between modality B and the output together drive the pre-fusion layer weights of modality B to escape from the zero fixed point. Thus having more shared post-fusion layers makes learning one modality more helpful for learning the other, shortening the unimodal phase. In essence, an early fusion point allows the weaker modality to benefit from the stronger modality’s learning in the post-fusion layers.

We also note that the initialization scale affects the time ratio in intermediate fusion networks as demonstrated in Fig. 5d. Even amongst cases that all fall into the rich feature learning regime, the initialization scale has an effect on the time ratio, with a larger time ratio for a larger initialization scale. In Fig. 5d, the simulations (circles) slightly deviate from theoretical predictions (lines) because our theoretical prediction is derived with small initialization and thus is less accurate for larger initialization. Nonetheless, the monotonic trend is well captured.

In summary, a deeper fusion layer, a larger input-output correlation ratio, stronger correlations between input modalities, and sometimes a smaller initialization scale all prolong the unimodal phase in the joint training of deep multimodal linear networks with small initialization.

5 DISCUSSION

We investigated the duration of the unimodal phase, mis-attribution, and superficial modality preference in deep intermediate and late fusion linear networks. We empirically find that our results, derived for linear networks, carry over to two-layer ReLU networks when the target task is linear, which aligns with the intuitions from a line of studies on two-layer ReLU networks (Sarussi et al., 2021; Phuong & Lampert, 2021; Min et al., 2023; Timor et al., 2023). We simulate and present the duration of the unimodal phase and the amount of mis-attribution in two-layer late fusion ReLU networks in Figs. 3b and 3c (crosses), which closely follow the theoretical predictions and results of their linear counterparts. The loss and weights trajectories shown in Fig. 10 are also qualitatively the same as the trajectories in two-layer linear networks in Fig. 2, except that learning is about two times slower and the converged total weights are two times larger.

In practice, multimodal data is correlated and heterogeneous. We explained the unimodal bias induced by correlation in deep linear networks. More generally, the heterogeneity in data and nonlinearity in neural networks can induce behaviors that do not arise in linear networks and linear tasks, which are out of the scope of this paper. Nonetheless, we present a simple nonlinear example in the hope of inspiring future work. Consider learning $y = \mathbf{x}_A + \text{XOR}(\mathbf{x}_B)$, where $\mathbf{x}_A \in \mathbb{R}$, $\mathbf{x}_B \in \{[1, 1], [1, -1], [-1, 1], [-1, -1]\}$ and $\text{XOR}(\mathbf{x}_B)$ refers to performing XOR to the two dimensions of \mathbf{x}_B . We observe that two-layer late fusion ReLU networks always learn this task successfully, forming the four perpendicular XOR features perfectly as shown in Figs. 11b, 11d and 11f. However, two-layer early fusion ReLU networks do not learn consistent XOR features and can even fail to learn this task when the variance of \mathbf{x}_A is large so that exploiting the linear modality causes fatal perturbations to extracting features from the XOR modality as shown in Figs. 11a, 11c and 11e. In this nonlinear example, late fusion networks are advantageous in terms of extracting heterogeneous features from each input modality.

We hypothesize that the practical choice of the fusion layer depth is a trade-off between alleviating unimodal bias and learning unimodal features. A shallower fusion layer helps alleviate unimodal bias because modalities can cooperate reciprocally to learn the synergistic computation. Meanwhile, a deeper fusion layer helps unimodal feature learning because the network can operate more independently to learn the unique computation of extracting heterogeneous features from each modality. We hope our work contributes to a better understanding of this tradeoff, ultimately leading to more systematic architectural choices and improved multimodal learning algorithms.

REFERENCES

- Madhu S. Advani, Andrew M. Saxe, and Haim Sompolinsky. High-dimensional dynamics of generalization error in neural networks. *Neural Networks*, 132:428–446, 2020. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2020.08.022>. URL <https://www.sciencedirect.com/science/article/pii/S0893608020303117>.
- Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1955–1960, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1203. URL <https://aclanthology.org/D16-1203>.
- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. Don’t just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 244–253. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/arora18a.html>.
- Alexander Atanasov, Blake Bordelon, and Cengiz Pehlevan. Neural networks as kernel learners: The silent alignment effect. In *International Conference on Learning Representations, 2022*. URL <https://openreview.net/forum?id=1NvflqAdoom>.
- Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2): 423–443, 2019. doi: 10.1109/TPAMI.2018.2798607.
- George Barnum, Sabera J Talukder, and Yisong Yue. On the benefits of early fusion in multimodal representation learning. In *NeurIPS 2020 Workshop SVRHM*, 2020. URL <https://openreview.net/forum?id=FMJ5e0IoFFk>.
- Remi Cadene, Corentin Dancette, Hedi Ben younes, Matthieu Cord, and Devi Parikh. Rubi: Reducing unimodal biases for visual question answering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/51d92belc60d1db1d2e5e7a07da55b26-Paper.pdf.
- Chenzhuang Du, Jiaye Teng, Tingle Li, Yichen Liu, Tianyuan Yuan, Yue Wang, Yang Yuan, and Hang Zhao. On uni-modal feature learning in supervised multi-modal learning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 8632–8656. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/du23e.html>.
- Simon S Du, Wei Hu, and Jason D Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/fe131d7f5a6b38b23cc967316c13dae2-Paper.pdf.
- Kenji Fukumizu. Effect of batch learning in multilayer neural networks. *Gen*, 1(04):1E–03, 1998.
- Konrad Gadzicki, Razieh Khamsehashari, and Christoph Zetsche. Early vs late fusion in multimodal convolutional neural networks. In *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, pp. 1–6, 2020. doi: 10.23919/FUSION45008.2020.9190246.

- Itai Gat, Idan Schwartz, and Alex Schwing. Perceptual score: What data modalities does your model perceive? In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 21630–21643. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/b51a15f382ac914391a58850ab343b00-Paper.pdf.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Jack Hessel and Lillian Lee. Does my multimodal model learn cross-modal interactions? it’s harder to tell than you might think! In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 861–877, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.62. URL <https://aclanthology.org/2020.emnlp-main.62>.
- Yu Huang, Chenzhuang Du, Zihui Xue, Xuanyao Chen, Hang Zhao, and Longbo Huang. What makes multi-modal learning better than single (provably). In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 10944–10956. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/5aa3405a3f865c10f420a4a7b55cbff3-Paper.pdf.
- Yu Huang, Junyang Lin, Chang Zhou, Hongxia Yang, and Longbo Huang. Modality competition: What makes joint training of multi-modal network fail in deep learning? (Provably). In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 9226–9259. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/huang22e.html>.
- Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- Arthur Jacot, François Ged, Berfin Şimşek, Clément Hongler, and Franck Gabriel. Saddle-to-saddle dynamics in deep linear networks: Small initialization training, symmetry, and sparsity. *arXiv preprint arXiv:2106.15933*, 2021.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HJflg30qKX>.
- Michael Kleinman, Alessandro Achille, and Stefano Soatto. Critical learning periods for multisensory integration in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24296–24305, June 2023.
- Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations and recent trends in multimodal machine learning: Principles, challenges, and open questions. *arXiv preprint arXiv:2209.03430*, 2022.
- Taro Makino, Yixin Wang, Krzysztof J. Geras, and Kyunghyun Cho. Detecting incidental correlation in multimodal learning via latent variable modeling. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=QoRo9QmOAr>.
- Hancheng Min, René Vidal, and Enrique Mallada. Early neuron alignment in two-layer relu networks with small initialization. *arXiv preprint arXiv:2307.12851*, 2023.
- Xiaokang Peng, Yake Wei, Andong Deng, Dong Wang, and Di Hu. Balanced multimodal learning via on-the-fly gradient modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8238–8247, June 2022.
- Scott Pesme and Nicolas Flammarion. Saddle-to-saddle dynamics in diagonal linear networks. *arXiv preprint arXiv:2304.00488*, 2023.

- Mary Phuong and Christoph H Lampert. The inductive bias of relu networks on orthogonally separable data. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=krz7T0xU9Z_.
- Dhanesh Ramachandram and Graham W. Taylor. Deep multimodal learning: A survey on recent advances and trends. *IEEE Signal Processing Magazine*, 34(6):96–108, 2017. doi: 10.1109/MSP.2017.2738401.
- Roei Sarussi, Alon Brutzkus, and Amir Globerson. Towards understanding learning in neural networks with linear teachers. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9313–9322. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/sarussi21a.html>.
- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations*, 2014. URL <https://arxiv.org/abs/1312.6120>.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23):11537–11546, 2019. doi: 10.1073/pnas.1820226116. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1820226116>.
- Jianghong Shi, Eric Shea-Brown, and Michael Buice. Learning dynamics of deep linear networks with multiple pathways. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 34064–34076. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/dc3ca8bcd613e43ce540352b58d55d6d-Paper-Conference.pdf.
- Nadav Timor, Gal Vardi, and Ohad Shamir. Implicit regularization towards rank minimization in relu networks. In Shipra Agrawal and Francesco Orabona (eds.), *Proceedings of The 34th International Conference on Algorithmic Learning Theory*, volume 201 of *Proceedings of Machine Learning Research*, pp. 1429–1459. PMLR, 20 Feb–23 Feb 2023.
- Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks hard? In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12692–12702, 2020. doi: 10.1109/CVPR42600.2020.01271.
- Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 24043–24055. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/wu22d.html>.
- Bo Xie, Yingyu Liang, and Le Song. Diverse Neural Network Learns True Target Functions. In Aarti Singh and Jerry Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1216–1224. PMLR, 20–22 Apr 2017. URL <https://proceedings.mlr.press/v54/xie17a.html>.

APPENDIX

A GRADIENT DESCENT IN DEEP MULTIMODAL LINEAR NETWORKS

We derive the full-batch gradient descent dynamics in deep multimodal linear networks with learning rate η . In pre-fusion layers $1 \leq l \leq L_f$, the gradient update is

$$\begin{aligned} \Delta \mathbf{W}_A^l &= -\eta \frac{\partial \mathcal{L}}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial \mathbf{W}_A^l} \\ &= \eta \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_A^i \right)^\top (y - \hat{y}) \left(\prod_{i=1}^{l-1} \mathbf{W}_A^i \mathbf{x}_A \right)^\top \\ &= \eta \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_A^i \right)^\top (y - \mathbf{W}_A^{\text{tot}} \mathbf{x}_A - \mathbf{W}_B^{\text{tot}} \mathbf{x}_B) \mathbf{x}_A^\top \left(\prod_{i=1}^{l-1} \mathbf{W}_A^i \right)^\top \\ &= \eta \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_A^i \right)^\top (\Sigma_{y\mathbf{x}_A} - \mathbf{W}_A^{\text{tot}} \Sigma_A - \mathbf{W}_B^{\text{tot}} \Sigma_{BA}) \left(\prod_{i=1}^{l-1} \mathbf{W}_A^i \right)^\top \end{aligned} \quad (13a)$$

$$\Delta \mathbf{W}_B^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_B^i \right)^\top (\Sigma_{y\mathbf{x}_B} - \mathbf{W}_A^{\text{tot}} \Sigma_{BA} - \mathbf{W}_B^{\text{tot}} \Sigma_B) \left(\prod_{i=1}^{l-1} \mathbf{W}_B^i \right)^\top \quad (13b)$$

In post-fusion layers $L_f + 1 \leq l \leq L$,

$$\begin{aligned} \Delta \mathbf{W}^l &= -\eta \frac{\partial \mathcal{L}}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial \mathbf{W}^l} \\ &= \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (y - \hat{y}) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_A^i \mathbf{x}_A \right)^\top \\ &\quad + \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (y - \hat{y}) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_B^i \mathbf{x}_B \right)^\top \\ &= \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (y - \mathbf{W}_A^{\text{tot}} \mathbf{x}_A - \mathbf{W}_B^{\text{tot}} \mathbf{x}_B) \mathbf{x}_A^\top \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_A^i \right)^\top \\ &\quad + \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (y - \mathbf{W}_A^{\text{tot}} \mathbf{x}_A - \mathbf{W}_B^{\text{tot}} \mathbf{x}_B) \mathbf{x}_B^\top \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_B^i \right)^\top \\ &= \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (\Sigma_{y\mathbf{x}_A} - \mathbf{W}_A^{\text{tot}} \Sigma_A - \mathbf{W}_B^{\text{tot}} \Sigma_{BA}) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_A^i \mathbf{x}_A \right)^\top \\ &\quad + \eta \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top (\Sigma_{y\mathbf{x}_B} - \mathbf{W}_A^{\text{tot}} \Sigma_{AB} - \mathbf{W}_B^{\text{tot}} \Sigma_B) \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_B^i \mathbf{x}_B \right)^\top \end{aligned} \quad (14)$$

In the limit of small learning rate, the difference equations Eqs. (13) and (14) are well approximated by the differential equations Eqs. (3) and (4) in the main text.

B TWO-LAYER EARLY FUSION LINEAR NETWORK FIXED POINTS

A two-layer early fusion linear network is described as $\hat{y} = \mathbf{W}^2 \mathbf{W}^1 \mathbf{x}$. The gradient descent dynamics are

$$\tau \dot{\mathbf{W}}^1 = \mathbf{W}^{2\top} (\boldsymbol{\Sigma}_{y\mathbf{x}} - \mathbf{W}^2 \mathbf{W}^1 \boldsymbol{\Sigma}), \quad \tau \dot{\mathbf{W}}^2 = (\boldsymbol{\Sigma}_{y\mathbf{x}} - \mathbf{W}^2 \mathbf{W}^1 \boldsymbol{\Sigma}) \mathbf{W}^{1\top}. \quad (15)$$

By setting to the gradients to zero, we find that there are two manifolds of fixed points:

$$\mathbf{W}^2 = \mathbf{0}, \mathbf{W}^1 = \mathbf{0} \quad \Rightarrow \quad \mathcal{M}_0 = \{\mathbf{W} | \mathbf{W} = \mathbf{0}\}, \quad (16a)$$

$$\boldsymbol{\Sigma}_{y\mathbf{x}} - \mathbf{W}^2 \mathbf{W}^1 \boldsymbol{\Sigma} = \mathbf{0} \quad \Rightarrow \quad \mathcal{M}_* = \{\mathbf{W} | \mathbf{W}^{\text{tot}} = \boldsymbol{\Sigma}_{y\mathbf{x}} \boldsymbol{\Sigma}^{-1}\}. \quad (16b)$$

C TWO-LAYER LATE FUSION LINEAR NETWORK DYNAMICS

A two-layer late fusion linear network is described as $\hat{y} = \mathbf{W}_A^2 \mathbf{W}_A^1 \mathbf{x}_A + \mathbf{W}_B^2 \mathbf{W}_B^1 \mathbf{x}_B$. The gradient descent dynamics are

$$\tau \dot{\mathbf{W}}_A^1 = \mathbf{W}_A^{2\top} (\boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_A - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_{BA}), \quad (17a)$$

$$\tau \dot{\mathbf{W}}_A^2 = (\boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_A - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_{BA}) \mathbf{W}_A^{1\top}, \quad (17b)$$

$$\tau \dot{\mathbf{W}}_B^1 = \mathbf{W}_B^{2\top} (\boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_{AB} - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_B), \quad (17c)$$

$$\tau \dot{\mathbf{W}}_B^2 = (\boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_{AB} - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_B) \mathbf{W}_B^{1\top}. \quad (17d)$$

C.1 FIXED POINTS AND SADDLES

By setting the gradients in Eq. (17) to zero, we find that the two manifolds of fixed points in Eq. (5) exist in two-layer late fusion linear networks as well.

$$\begin{cases} \mathbf{W}_A^2 = \mathbf{0}, \mathbf{W}_A^1 = \mathbf{0} \\ \mathbf{W}_B^2 = \mathbf{0}, \mathbf{W}_B^1 = \mathbf{0} \end{cases} \Rightarrow \mathcal{M}_0 = \{\mathbf{W} | \mathbf{W} = \mathbf{0}\}, \quad (18a)$$

$$\begin{cases} \boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_A - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_{BA} = \mathbf{0} \\ \boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_{AB} - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_B = \mathbf{0} \end{cases} \Rightarrow \mathcal{M}_* = \{\mathbf{W} | \mathbf{W}^{\text{tot}} = \boldsymbol{\Sigma}_{y\mathbf{x}} \boldsymbol{\Sigma}^{-1}\}. \quad (18b)$$

In addition, there are two manifolds of saddles:

$$\begin{cases} \boldsymbol{\Sigma}_{y\mathbf{x}_A} - \mathbf{W}_A^2 \mathbf{W}_A^1 \boldsymbol{\Sigma}_A = \mathbf{0} \\ \mathbf{W}_B^2 = \mathbf{0}, \mathbf{W}_B^1 = \mathbf{0} \end{cases} \Rightarrow \mathcal{M}_A = \{\mathbf{W} | \mathbf{W}_A^{\text{tot}} = \boldsymbol{\Sigma}_{y\mathbf{x}_A} \boldsymbol{\Sigma}_A^{-1}, \mathbf{W}_B^{\text{tot}} = \mathbf{0}\}, \quad (19a)$$

$$\begin{cases} \mathbf{W}_A^2 = \mathbf{0}, \mathbf{W}_A^1 = \mathbf{0} \\ \boldsymbol{\Sigma}_{y\mathbf{x}_B} - \mathbf{W}_B^2 \mathbf{W}_B^1 \boldsymbol{\Sigma}_B = \mathbf{0} \end{cases} \Rightarrow \mathcal{M}_B = \{\mathbf{W} | \mathbf{W}_A^{\text{tot}} = \mathbf{0}, \mathbf{W}_B^{\text{tot}} = \boldsymbol{\Sigma}_{y\mathbf{x}_B} \boldsymbol{\Sigma}_B^{-1}\}. \quad (19b)$$

We plot the four manifolds $\mathcal{M}_0, \mathcal{M}_*, \mathcal{M}_A, \mathcal{M}_B$ in Fig. 6 in a scalar case introduced in Fig. 2d.

C.2 A SOLVABLE SIMPLE CASE

If the two modalities are not correlated with each other ($\boldsymbol{\Sigma}_{AB} = \mathbf{0}$) and have white correlations ($\boldsymbol{\Sigma}_A = \sigma_A^2 \mathbf{I}, \boldsymbol{\Sigma}_B = \sigma_B^2 \mathbf{I}$) the dynamics is equivalent to two separately trained unimodal two-layer linear networks with whitened input, whose solution has been derived in Saxe et al. (2014). The time course solution of total weights is

$$\mathbf{W}_A^{\text{tot}}(t) = u_A(t) \sigma_A^{-2} \boldsymbol{\Sigma}_{y\mathbf{x}_A}, \quad u_A(t) = \left[\left(\frac{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|}{\sigma_A^2 u_A(0)} - 1 \right) e^{-2\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\| \frac{t}{\tau}} + 1 \right]^{-1}, \quad (20a)$$

$$\mathbf{W}_B^{\text{tot}}(t) = u_B(t) \sigma_B^{-2} \boldsymbol{\Sigma}_{y\mathbf{x}_B}, \quad u_B(t) = \left[\left(\frac{\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|}{\sigma_B^2 u_B(0)} - 1 \right) e^{-2\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\| \frac{t}{\tau}} + 1 \right]^{-1}. \quad (20b)$$

The total weights for both modalities only evolve in scale along the pseudo-inverse solution direction. The scale variables $u_A(t), u_B(t)$ go through the same sigmoidal growth while modality A grows approximately $\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|/\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|$ times faster.




Figure 6: Phase portrait corresponding to learning dynamics in Fig. 2d. The horizontal axis u_A is the Frobenius norm of a layer in branch A as defined in Eq. (22); and similarly for the vertical axis u_B . Two-layer late fusion linear networks initialized with small weights either go along the solid or the dotted green path to reach the solid red dot stable fixed point in \mathcal{M}_* at convergence. Via either path, the network visits and lingers in a saddle in manifold \mathcal{M}_A , giving rise to the unimodal phase.

C.3 GENERAL CASES

In the general case of having arbitrary correlation matrices, the analytical solution in Eq. (20) does not hold since the total weights not only grow in one fixed direction but also rotate. We then focus on the early phase dynamics to compute the duration of the unimodal phase.

C.3.1 WHICH MODALITY IS LEARNED FIRST AND WHEN?

We assume that the modality learned first is modality A and specify the time t_A and the condition for prioritizing modality A in the following. During time 0 to t_A , the total weights of both modalities have not started to grow appreciably. Hence, the dynamics during time $0 \sim t_A$ in leading order approximation are

$$\tau \dot{\mathbf{W}}_A^1 = \mathbf{W}_A^{2\top} \Sigma_{y\mathbf{x}_A}, \quad \tau \dot{\mathbf{W}}_A^2 = \Sigma_{y\mathbf{x}_A} \mathbf{W}_A^{1\top}; \quad (21a)$$

$$\tau \dot{\mathbf{W}}_B^1 = \mathbf{W}_B^{2\top} \Sigma_{y\mathbf{x}_B}, \quad \tau \dot{\mathbf{W}}_B^2 = \Sigma_{y\mathbf{x}_B} \mathbf{W}_B^{1\top}. \quad (21b)$$

The first-layer weights $\mathbf{W}_A^1, \mathbf{W}_B^1$ align with $\Sigma_{y\mathbf{x}_A}, \Sigma_{y\mathbf{x}_B}$ respectively in the early phase when their scale has not grown appreciably (Atanasov et al., 2022). We further have the balancing property (Du et al., 2018; Ji & Telgarsky, 2019) between the two layers of modality A

$$\mathbf{W}_A^1 \mathbf{W}_A^{1\top} = \mathbf{W}_A^{2\top} \mathbf{W}_A^2 \Rightarrow \|\mathbf{W}_A^1\|_F = \|\mathbf{W}_A^2\| \stackrel{\text{def}}{=} u_A. \quad (22)$$

We conduct the following change of variable

$$\mathbf{W}_A^1 = u_A(t) \mathbf{r}_A^1 \frac{\Sigma_{y\mathbf{x}_A}}{\|\Sigma_{y\mathbf{x}_A}\|}, \quad \mathbf{W}_A^2 = u_A(t) \mathbf{r}_A^{1\top}, \quad (23)$$

where \mathbf{r}_A^1 is a fixed unit norm column vector representing the freedom in the hidden layer and u_A is a scalar representing the norm of the two balancing layers (Advani et al., 2020). Substituting the variables, we can re-write the dynamics of \mathbf{W}_A^2 and obtain an ordinary differential equation about u_A :

$$\tau \dot{u}_A \mathbf{r}_A^{1\top} = \Sigma_{y\mathbf{x}_A} \frac{\Sigma_{y\mathbf{x}_A}^\top}{\|\Sigma_{y\mathbf{x}_A}\|} \mathbf{r}_A^1 u_A \Rightarrow \tau \dot{u}_A = \|\Sigma_{y\mathbf{x}_A}\| u_A. \quad (24)$$

Separating variables and integrating both sides, we get

$$t = \tau \|\Sigma_{y\mathbf{x}_A}\|^{-1} (\ln u_A - \ln u_0), \quad (25)$$

where u_0 denotes $u_0 = u_A(0)$. Because the initialization u_0 is very small compared to $\|\Sigma_{y\mathbf{x}_A}\|$, the time that u_A grows to be comparable with $\|\Sigma_{y\mathbf{x}_A}\|$ is

$$t_A \approx \tau \|\Sigma_{y\mathbf{x}_A}\|^{-1} \ln \frac{1}{u_0}, \quad (26)$$

From Eq. (26), we can see that the condition for modality A to be learned first is that $\|\Sigma_{yx_A}\| > \|\Sigma_{yx_B}\|$.

$$\tau\|\Sigma_{yx_A}\|^{-1} \ln \frac{1}{u_0} < \tau\|\Sigma_{yx_B}\|^{-1} \ln \frac{1}{u_0} \Leftrightarrow \|\Sigma_{yx_A}\| > \|\Sigma_{yx_B}\|. \quad (27)$$

C.3.2 WHEN IS THE SECOND MODALITY LEARNED?

We next calculate the time t_B when modality B is learned. During time 0 to t_A , the weight dynamics of modality B takes the same form as modality A as described in Eq. (21b). Because the layers in the modality B branch are balanced and \mathbf{W}_B^1 aligns with the rank-one direction Σ_{yx_B} , we again change variables and re-write the dynamics of \mathbf{W}_B^2 , obtaining an ordinary differential equation about the norm of the two balanced layers $u_B = \|\mathbf{W}_B^1\|_F = \|\mathbf{W}_B^2\|$:

$$\tau \dot{u}_B = \|\Sigma_{yx_B}\| u_B \Rightarrow t = \tau \|\Sigma_{yx_B}\|^{-1} (\ln u_B - \ln u_0), t \in [0, t_A]. \quad (28)$$

We assume $u_B(0) = u_0$, since the initialization is small and the number of hidden neurons in the modality B branch is of the same order as modality A. We plug t_A from Eq. (26) into Eq. (28) and get

$$\ln u_B(t_A) = (1 - \|\Sigma_{yx_A}\|^{-1} \|\Sigma_{yx_B}\|) \ln u_0. \quad (29)$$

We then look into the dynamics during the time lag from t_A to t_B . During time t_A to t_B , the weights of modality B are still small and negligible to leading ordering approximation. Meanwhile, the weights of modality A have grown to be $\mathbf{W}_A^{\text{tot}} = \Sigma_{yx_A} \Sigma_A^{-1}$, which changes the dynamics of modality B when the cross-covariance $\Sigma_{AB} \neq \mathbf{0}$. Taking this into consideration, the dynamics of modality B during t_A to t_B is

$$\tau \dot{\mathbf{W}}_B^1 = \mathbf{W}_B^{2\top} \tilde{\Sigma}_{yx_B}, \tau \dot{\mathbf{W}}_B^2 = \tilde{\Sigma}_{yx_B} \mathbf{W}_B^{1\top}, t \in (t_A, t_B), \quad (30)$$

where $\tilde{\Sigma}_{yx_B} = \Sigma_{yx_B} - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB}$. The first-layer weights \mathbf{W}_B^1 rapidly rotate from Σ_{yx_B} to $\tilde{\Sigma}_{yx_B}$ at time t_A and continue to evolve along $\tilde{\Sigma}_{yx_B}$ during t_A to t_B . Through the same manner of changing variables, we obtain the ordinary differential equation about u_B during t_A to t_B :

$$\tau \dot{u}_B = \|\tilde{\Sigma}_{yx_B}\| u_B \Rightarrow t - t_A = \tau \|\tilde{\Sigma}_{yx_B}\|^{-1} (\ln u_B - \ln u_B(t_A)), t \in (t_A, t_B). \quad (31)$$

Plugging in $u_B(t_A)$ obtained in Eq. (29), we get the time when u_B grows to be comparable with $\|\Sigma_{yx_B}\|$:

$$t_B \approx t_A - \tau \|\tilde{\Sigma}_{yx_B}\|^{-1} \ln u_B(t_A) \approx t_A + \tau \frac{1 - \|\Sigma_{yx_A}\|^{-1} \|\Sigma_{yx_B}\|}{\|\tilde{\Sigma}_{yx_B}\|} \ln \frac{1}{u_0}. \quad (32)$$

Dividing Eq. (32) by Eq. (26), we obtain the time ratio Eq. (8) in the main text:

$$\frac{t_B}{t_A} = 1 + \frac{\|\Sigma_{yx_A}\| - \|\Sigma_{yx_B}\|}{\|\Sigma_{yx_B} - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{AB}\|}. \quad (33)$$




Figure 7: Loss and weight trajectories in a two-layer late fusion network when modalities have collinearity. Inputs \mathbf{x}_A and \mathbf{x}_B are scalars with covariance matrix $\Sigma = [4, 2; 2, 1]$. The target output is generated as $y = \mathbf{x}_A + \mathbf{x}_B$.

Note that if \mathbf{x}_A and \mathbf{x}_B have collinearity, the denominator in Eq. (33) is zero and the time ratio is ∞ . As shown in Fig. 7, the late fusion network learns to fit the output only with modality A and modality B will never be learned.

D DEEP INTERMEDIATE & LATE FUSION LINEAR NETWORKS

D.1 BALANCING PROPERTIES

The full dynamics of deep intermediate and late fusion linear networks are given in Eqs. (3) and (4). Between pre-fusion layers and between post-fusion layers, the standard balancing property (Du et al., 2018; Ji & Telgarsky, 2019) holds true:

$$\mathbf{W}_A^1 \mathbf{W}_A^{1\top} = \mathbf{W}_A^2 \mathbf{W}_A^{2\top} = \dots = \mathbf{W}_A^{L_f} \mathbf{W}_A^{L_f\top}, \quad (34a)$$

$$\mathbf{W}_B^1 \mathbf{W}_B^{1\top} = \mathbf{W}_B^2 \mathbf{W}_B^{2\top} = \dots = \mathbf{W}_B^{L_f} \mathbf{W}_B^{L_f\top}, \quad (34b)$$

$$\mathbf{W}^{L_f+1} \mathbf{W}^{L_f+1\top} = \mathbf{W}^{L_f+2} \mathbf{W}^{L_f+2\top} = \dots = \mathbf{W}^L \mathbf{W}^L. \quad (34c)$$

Between a pre-fusion layer and a post-fusion layer, the balancing condition takes a slightly different form:

$$\mathbf{W}_A^{L_f} \mathbf{W}_A^{L_f\top} + \mathbf{W}_B^{L_f} \mathbf{W}_B^{L_f\top} = \mathbf{W}^{L_f+1} \mathbf{W}^{L_f+1\top}. \quad (35)$$

Based on the standard balancing properties in Eq. (34), we have equal norm between pre-fusion layers and between post-fusion layers.

$$\|\mathbf{W}_A^1\|_F = \|\mathbf{W}_A^2\|_F = \dots = \|\mathbf{W}_A^{L_f}\|_F \stackrel{\text{def}}{=} u_A, \quad (36a)$$

$$\|\mathbf{W}_B^1\|_F = \|\mathbf{W}_B^2\|_F = \dots = \|\mathbf{W}_B^{L_f}\|_F \stackrel{\text{def}}{=} u_B, \quad (36b)$$

$$\|\mathbf{W}^{L_f+1}\|_F = \|\mathbf{W}^{L_f+2}\|_F = \dots = \|\mathbf{W}^L\|_F \stackrel{\text{def}}{=} u. \quad (36c)$$

Based on Eq. (35), we have the balancing property between the norm of a pre-fusion layer and a post-fusion layer.

$$u_A^2 + u_B^2 = u^2. \quad (37)$$

In addition to the balancing norm, we infer from the balancing properties that post-fusion layers have rank-one structures as in standard linear networks (Ji & Telgarsky, 2019; Atanasov et al., 2022). The pre-fusion layers are not necessarily rank-one due to the different balancing property in Eq. (35). However, guided by empirical observations, we make the ansatz that the weights in pre-fusion layers are also rank-one, which will enable us to conduct the change of variables similar to what we have done in Appendix C.3.

D.2 TIME RATIO FOR $L_f \neq 2$

D.2.1 WHICH MODALITY IS LEARNED FIRST AND WHEN?

We adopt the convention that the modality learned first is modality A. During time 0 to t_A , the dynamics in leading order approximation are

$$\tau \dot{\mathbf{W}}_A^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_A^i \right)^\top \Sigma_{y\mathbf{x}_A} \left(\prod_{i=1}^{l-1} \mathbf{W}_A^i \right)^\top, \quad 1 \leq l \leq L_f, \quad (38a)$$

$$\tau \dot{\mathbf{W}}_B^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_B^i \right)^\top \Sigma_{y\mathbf{x}_B} \left(\prod_{i=1}^{l-1} \mathbf{W}_B^i \right)^\top, \quad 1 \leq l \leq L_f, \quad (38b)$$

$$\begin{aligned} \tau \dot{\mathbf{W}}^l &= \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top \Sigma_{y\mathbf{x}_A} \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_A^i \right)^\top \\ &+ \left(\prod_{j=l+1}^L \mathbf{W}^j \right)^\top \Sigma_{y\mathbf{x}_B} \left(\prod_{j=L_f+1}^{l-1} \mathbf{W}^j \prod_{i=1}^{L_f} \mathbf{W}_B^i \right)^\top, \quad L_f + 1 \leq l \leq L. \end{aligned} \quad (38c)$$

We make the following change of variables

$$\mathbf{W}_A^1 = u_A(t) \mathbf{r}_A^1 \frac{\boldsymbol{\Sigma}_{y\mathbf{x}_A}}{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|}, \quad \mathbf{W}_B^1 = u_B(t) \mathbf{r}_B^1 \frac{\boldsymbol{\Sigma}_{y\mathbf{x}_B}}{\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|}; \quad (39a)$$

$$\mathbf{W}_A^i = u_A(t) \mathbf{r}_A^i \mathbf{r}_A^{i-1\top}, \quad \mathbf{W}_B^i = u_B(t) \mathbf{r}_B^i \mathbf{r}_B^{i-1\top}, \quad 2 \leq i \leq L_f - 1; \quad (39b)$$

$$\mathbf{W}_A^{L_f} = u_A(t) \mathbf{r}_A^{L_f} \mathbf{r}_A^{L_f-1\top}, \quad \mathbf{W}_B^{L_f} = u_B(t) \mathbf{r}_B^{L_f} \mathbf{r}_B^{L_f-1\top}; \quad (39c)$$

$$\mathbf{W}^j = u(t) \mathbf{r}^j \mathbf{r}^{j-1\top}, \quad L_f + 1 \leq j \leq L - 1; \quad (39d)$$

$$\mathbf{W}^L = u(t) \mathbf{r}^L \mathbf{r}^{L\top}. \quad (39e)$$

where all \mathbf{r} are fixed unit norm column vectors representing the freedom in hidden layers. By substituting Eq. (39) into Eq. (38), we reduce the dynamics during time 0 to t_A to a three-dimensional dynamical system about u_A , u_B and u :

$$\tau \dot{u}_A = \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\| u_A^{L_f-1} u^{L-L_f}, \quad (40a)$$

$$\tau \dot{u}_B = \|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\| u_B^{L_f-1} u^{L-L_f}, \quad (40b)$$

$$\tau \dot{u} = \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\| u_A^{L_f} u^{L-L_f-1} + \|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\| u_B^{L_f} u^{L-L_f-1}. \quad (40c)$$

We divide Eq. (40b) by Eq. (40a) and reveal an equality between the two branches of a pre-fusion layer:

$$\frac{du_B}{du_A} = \frac{\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\| u_B^{L_f-1}}{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\| u_A^{L_f-1}} \Rightarrow \frac{u_B^{2-L_f} - u_0^{2-L_f}}{u_A^{2-L_f} - u_0^{2-L_f}} = \frac{\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|}{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|} \stackrel{\text{def}}{=} k, \quad L_f \neq 2, \quad (41)$$

We assume $L_f \neq 2$ in this section and handle the case of $L_f = 2$ in Appendix D.3. By utilizing two equality properties in Eq. (41) and Eq. (37), we reduce Eq. (40a) to an ordinary differential equation:

$$\tau \dot{u}_A = \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\| u_A^{L_f-1} \left[u_A^2 + \left(k u_A^{2-L_f} + (1-k) u_0^{2-L_f} \right)^{\frac{2}{2-L_f}} \right]^{\frac{L-L_f}{2}}. \quad (42)$$

The ordinary differential equation Eq. (42) is separable despite being cumbersome. By separating variables and integrating both sides, we can write time t as a function of u_A :

$$t = \int_{u_0}^{u_A} x^{1-L_f} \left[x^2 + \left(k x^{2-L_f} + (1-k) u_0^{2-L_f} \right)^{\frac{2}{2-L_f}} \right]^{\frac{L_f-L}{2}} dx. \quad (43)$$

The time when u_A grows to be comparable with $\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|$, for instance $\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|/2$, is

$$\begin{aligned} t_A &\approx \tau \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|^{-1} \int_{u_0}^{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|/2} u_A^{1-L_f} \left[u_A^2 + \left(k u_A^{2-L_f} + (1-k) u_0^{2-L_f} \right)^{\frac{2}{2-L_f}} \right]^{\frac{L_f-L}{2}} du_A \\ &\approx \tau \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|^{-1} u_0^{2-L} \int_1^\infty x^{1-L} \left[1 + \left(k + (1-k) x^{L_f-2} \right)^{\frac{2}{2-L_f}} \right]^{\frac{L_f-L}{2}} dx. \end{aligned} \quad (44)$$

From Eq. (44), we find that the condition for modality A to be learned first in deep intermediate and late fusion linear networks is the same as that for two-layer late fusion linear networks:

$$\tau \|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|^{-1} u_0^{2-L_f} I(L, L_f) < \tau \|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|^{-1} u_0^{2-L_f} I(L, L_f) \Leftrightarrow \frac{\|\boldsymbol{\Sigma}_{y\mathbf{x}_B}\|}{\|\boldsymbol{\Sigma}_{y\mathbf{x}_A}\|} \equiv k \in (0, 1), \quad (45)$$

where we use $I(L, L_f)$ to denote the integral

$$I(L, L_f) = \int_1^\infty x^{1-L} \left[1 + \left(k + (1-k) x^{L_f-2} \right)^{\frac{2}{2-L_f}} \right]^{\frac{L_f-L}{2}} dx. \quad (46)$$

D.2.2 WHEN IS THE SECOND MODALITY LEARNED?

We next compute the time t_B when modality B is learned. During time t_A to t_B , the network is in manifold \mathcal{M}_A where

$$\mathbf{W}_A^{\text{tot}} = \Sigma_{y\mathbf{x}_A} \Sigma_A^{-1} \Rightarrow u_A = u = \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{\frac{1}{L}}, t \in (t_A, t_B). \quad (47)$$

We plug $u_A(t_A)$, which is very large compared to u_0 , into Eq. (41) and obtain $u_B(t_A)$

$$u_B(t_A) = \left[u_A(t_A)^{2-L_f} + (1-k)u_0^{2-L_f} \right]^{\frac{1}{2-L_f}} \approx (1-k)^{\frac{1}{2-L_f}} u_0. \quad (48)$$

We then look into the dynamics during the time lag t_A to t_B . Substituting $\mathbf{W}_A^{\text{tot}} = \Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}$ into Eq. (3b), we get

$$\tau \dot{\mathbf{W}}_B^l = \left(\prod_{j=L_f+1}^L \mathbf{W}^j \prod_{i=l+1}^{L_f} \mathbf{W}_B^i \right)^\top \tilde{\Sigma}_{y\mathbf{x}_B} \left(\prod_{i=1}^{l-1} \mathbf{W}_B^i \right)^\top, 1 \leq l \leq L_f, \quad (49)$$

where $\tilde{\Sigma}_{y\mathbf{x}_B} = \Sigma_{y\mathbf{x}_B} - \Sigma_{y\mathbf{x}_A} \Sigma_A^{-1} \Sigma_{AB}$. The first-layer weights \mathbf{W}_B^1 rapidly rotate from $\Sigma_{y\mathbf{x}_B}$ to $\tilde{\Sigma}_{y\mathbf{x}_B}$ at time t_A and continues to evolve along $\tilde{\Sigma}_{y\mathbf{x}_B}$ during t_A to t_B . Through the same manner of changing variables, we obtain an ordinary differential equation for u_B during t_A to t_B :

$$\begin{aligned} \tau \dot{u}_B &= \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{1-\frac{L_f}{L}} \|\tilde{\Sigma}_{y\mathbf{x}_B}\| u_B^{L_f-1} \\ \Rightarrow t - t_A &= \tau \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{\frac{L_f}{L}-1} \|\tilde{\Sigma}_{y\mathbf{x}_B}\|^{-1} \frac{u_B^{2-L_f} - u_B(t_A)^{2-L_f}}{2-L_f}, t \in (t_A, t_B). \end{aligned} \quad (50)$$

Plugging in $u_B(t_A)$ obtained in Eq. (48), we get the time t_B :

$$\begin{aligned} t_B &= t_A + \tau \frac{u_B(t_B)^{2-L_f} - u_B(t_A)^{2-L_f}}{(2-L_f) \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{1-\frac{L_f}{L}} \|\tilde{\Sigma}_{y\mathbf{x}_B}\|} \\ &\approx t_A + \tau \frac{-u_B(t_A)^{2-L_f}}{(2-L_f) \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{1-\frac{L_f}{L}} \|\tilde{\Sigma}_{y\mathbf{x}_B}\|} \\ &\approx t_A + \tau \frac{1 - \|\Sigma_{y\mathbf{x}_A}\|^{-1} \|\Sigma_{y\mathbf{x}_B}\|}{(L_f - 2) \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{1-\frac{L_f}{L}} \|\tilde{\Sigma}_{y\mathbf{x}_B}\|} u_0^{2-L_f}. \end{aligned} \quad (51)$$

Dividing Eq. (51) by Eq. (44), we arrive at the time ratio Eq. (11) in the main text. For intermediate and late fusion linear networks $2 < L_f \leq L$, the time ratio is

$$\frac{t_B}{t_A} = 1 + \frac{(\|\Sigma_{y\mathbf{x}_A}\| - \|\Sigma_{y\mathbf{x}_B}\|) u_0^{L-L_f}}{(L_f - 2) \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{1-\frac{L_f}{L}} \|\Sigma_{y\mathbf{x}_B} - \Sigma_{y\mathbf{x}_A} \Sigma_A^{-1} \Sigma_{AB}\|} I(L, L_f)^{-1}, \quad (52)$$

where the integral $I(L, L_f)$ has been defined in Eq. (46).

D.3 TIME RATIO FOR $L_f = 2$

When the fusion layer is the second layer $L_f = 2$, the equality in Eq. (41) takes following form:

$$\frac{du_A}{du_B} = \frac{\|\Sigma_{y\mathbf{x}_A}\| u_A}{\|\Sigma_{y\mathbf{x}_B}\| u_B} \Rightarrow \frac{\ln u_A - \ln u_0}{\ln u_B - \ln u_0} = \frac{\|\Sigma_{y\mathbf{x}_A}\|}{\|\Sigma_{y\mathbf{x}_B}\|}. \quad (53)$$

Consequently, fusion at the second layer is a special case with slightly different expressions for the times. We follow the same procedure as Appendix D.2 and obtain the times for $2 = L_f < L$:

$$t_A \approx \tau \frac{u_0^{2-L}}{\|\Sigma_{y\mathbf{x}_A}\|} \int_1^\infty x^{-1} (x^2 + x^{2k})^{1-\frac{L}{2}} dx, \quad (54a)$$

$$t_B \approx t_A + \tau \frac{1 - \|\Sigma_{y\mathbf{x}_A}\|^{-1} \|\Sigma_{y\mathbf{x}_B}\|}{\|\Sigma_{y\mathbf{x}_B} - \Sigma_{y\mathbf{x}_A} \Sigma_A^{-1} \Sigma_{AB}\|} \|\Sigma_{y\mathbf{x}_A} \Sigma_A^{-1}\|^{\frac{2}{L}-1} \ln \frac{1}{u_0}. \quad (54b)$$

The time ratio is

$$\frac{t_B}{t_A} = 1 + \frac{(\|\Sigma_{y_{x_A}}\| - \|\Sigma_{y_{x_B}}\|)u_0^{L-2} \ln \frac{1}{u_0}}{(L_f - 2) \|\Sigma_{y_{x_B}} - \Sigma_{y_{x_A}} \Sigma_A^{-1} \Sigma_{AB}\| \|\Sigma_{y_{x_A}} \Sigma_A^{-1}\|^{1-\frac{2}{L}}} I(L, 2)^{-1}, \quad (55)$$

where the integral is given by

$$I(L, 2) = \int_1^\infty x^{-1} (x^2 + x^{2k})^{1-\frac{L}{2}} dx. \quad (56)$$

D.4 TIME RATIO FOR UNEQUAL DEPTH

Our time ratio calculations can be carried out for intermediate fusion linear networks with unequal depth between modalities. Consider a intermediate fusion linear network with L_c post-fusion layers, L_A pre-fusion layers for the modality A branch, and L_B pre-fusion layers for the modality B branch. Assuming $L_A, L_B > 2$ and modality A is learned first, the time ratio is

$$\frac{t_B}{t_A} = 1 + \frac{\frac{u_0^{L_c+L_A-L_B}}{L_B-2} \|\Sigma_{y_{x_A}}\| - \frac{u_0^{L_c}}{L_A-2} \|\Sigma_{y_{x_B}}\|}{\|\Sigma_{y_{x_A}} \Sigma_A^{-1}\|^{\frac{L_c}{L_A+L_c}} \|\Sigma_{y_{x_B}} - \Sigma_{y_{x_A}} \Sigma_A^{-1} \Sigma_{AB}\|} I(L_A, L_B, L_c)^{-1}, \quad (57)$$

where the integral is given by

$$I(L_A, L_B, L_c) = \int_1^\infty x^{1-L_A} \left[x^2 + \left(\frac{(L_B - 2) \|\Sigma_{y_{x_B}}\|}{(L_A - 2) \|\Sigma_{y_{x_A}}\|} u_0^{L_B-L_A} (x^{2-L_A} - 1) + 1 \right)^{\frac{2}{2-L_B}} \right]^{-\frac{L_c}{2}} dx. \quad (58)$$

As a sanity check, setting $L_A = L_B$ gives us back the same expression as Eq. (11).

E FEATURE EVOLUTION IN MULTIMODAL LINEAR NETWORKS

We compare the feature evolution in two-layer early fusion and late fusion linear networks to gain a more detailed understanding of their different learning dynamics. As studied by Atanasov et al. (2022), features of linear networks lie in the first-layer weights. We thus plot the first-layer weights in multimodal linear networks at different times of training to inspect the feature evolution.

E.1 FEATURE EVOLUTION IN EARLY FUSION LINEAR NETWORKS

In early fusion linear networks with small initialization, the balancing property $\mathbf{W}^1 \mathbf{W}^{1\top} = \mathbf{W}^{2\top} \mathbf{W}^2$ holds true throughout training, which implies \mathbf{W}^1 is rank-one throughout training (Du et al., 2018; Ji & Telgarsky, 2019). Specifically, \mathbf{W}^1 initially aligns with $\Sigma_{y_{x}}$ during the plateau and eventually aligns with $\Sigma_{y_{x}} \Sigma^{-1}$ after \mathbf{W}^1 have grown in scale and rotated during the brief transitional period. We illustrate this process with Fig. 8b and videos in the supplementary material. In deep early fusion linear networks, \mathbf{W}^1 behaves qualitatively the same.

E.2 FEATURE EVOLUTION IN INTERMEDIATE AND LATE FUSION LINEAR NETWORKS

In intermediate and late fusion linear networks, the balancing property takes a different form as given in Eqs. (34) and (35). Thus the first-layer weights are not constrained to be rank-one. Specifically, \mathbf{W}_A^1 grows during the first transitional period while \mathbf{W}_B^1 remains close to small initialization. After a unimodal phase, \mathbf{W}_B^1 starts to grow during the second transitional period while \mathbf{W}_A^1 stays unchanged, shrinks in scale, or expands in scale depending on modality B has zero, positive, or negative correlation with modality A. We illustrate this process with Fig. 8d and videos in the supplementary material. In deep late and intermediate fusion linear networks, \mathbf{W}^1 behaves qualitatively the same as in two-layer late fusion linear networks.

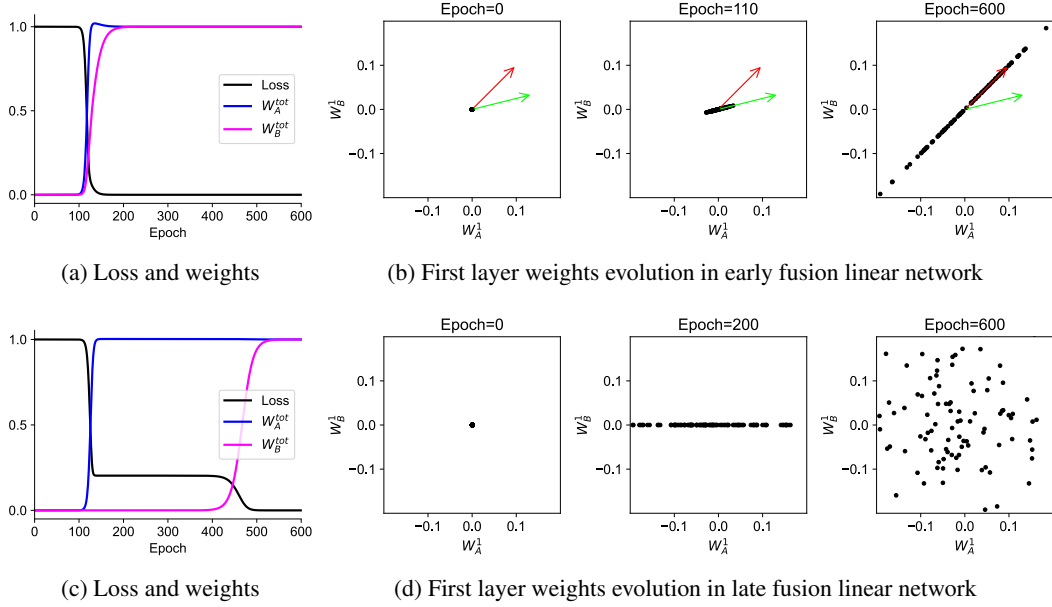


Figure 8: Feature evolution in two-layer early fusion and late fusion linear networks. We plot the feature evolution, which are first-layer weights in linear networks, corresponding to dynamics in Figs. 2b and 2d. (a) Time trajectories of loss and total weights of the same two-layer early fusion network as Fig. 2b. (b) First-layer weights at initialization, during training, and at convergence in (a). The green arrow denotes the direction of Σ_{yx} . The red arrow denotes the direction of $\Sigma_{yx}\Sigma^{-1}$. (c) Time trajectories of loss and total weights of the same two-layer late fusion network as Fig. 2d. (d) First-layer weights at initialization, during training, and at convergence in (c). Complementary videos can be found in the supplementary material.

F SUPERFICIAL MODALITY PREFERENCE

We impose that modality A contributes less to the output, meaning the loss would decrease more if the network visited a saddle in \mathcal{M}_B instead of \mathcal{M}_A as illustrated in Fig. 4a:

$$\mathcal{L}(\mathcal{M}_A) > \mathcal{L}(\mathcal{M}_B) \quad (59)$$

Substituting in the saddles defined in Eq. (6), we expand and simplify the two mean square losses:

$$\begin{aligned} \mathcal{L}(\mathcal{M}_A) &= \langle (y - \Sigma_{yx_A} \Sigma_A^{-1} x_A)^2 \rangle \\ &= \langle y^2 \rangle - 2 \Sigma_{yx_A} \Sigma_A^{-1} \langle y x_A \rangle + \Sigma_{yx_A} \Sigma_A^{-1} \langle x_A x_A^\top \rangle \Sigma_A^{-1} \Sigma_{yx_A}^\top \\ &= \langle y^2 \rangle - \Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{yx_A}^\top, \end{aligned} \quad (60a)$$

$$\begin{aligned} \mathcal{L}(\mathcal{M}_B) &= \langle (y - \Sigma_{yx_B} \Sigma_B^{-1} x_B)^2 \rangle \\ &= \langle y^2 \rangle - 2 \Sigma_{yx_B} \Sigma_B^{-1} \langle y x_B \rangle + \Sigma_{yx_B} \Sigma_B^{-1} \langle x_B x_B^\top \rangle \Sigma_B^{-1} \Sigma_{yx_B}^\top \\ &= \langle y^2 \rangle - \Sigma_{yx_B} \Sigma_B^{-1} \Sigma_{yx_B}^\top. \end{aligned} \quad (60b)$$

Plugging Eq. (60) into Eq. (59) gives us

$$\Sigma_{yx_A} \Sigma_A^{-1} \Sigma_{yx_A}^\top < \Sigma_{yx_B} \Sigma_B^{-1} \Sigma_{yx_B}^\top. \quad (61)$$

Since we are also assuming modality A is learned first, which is true when $\|\Sigma_{yx_A}\| > \|\Sigma_{yx_B}\|$. We thereby arrive at the two inequality conditions Eq. (9) in the main text.

G UNDERPARAMETERIZATION AND OVERPARAMETERIZATION IN MULTIMODAL LINEAR NETWORKS

We have studied multimodal networks trained with sufficient noiseless data so far. In this section, we consider the case where the training set contains noise and the amount of training samples vary.

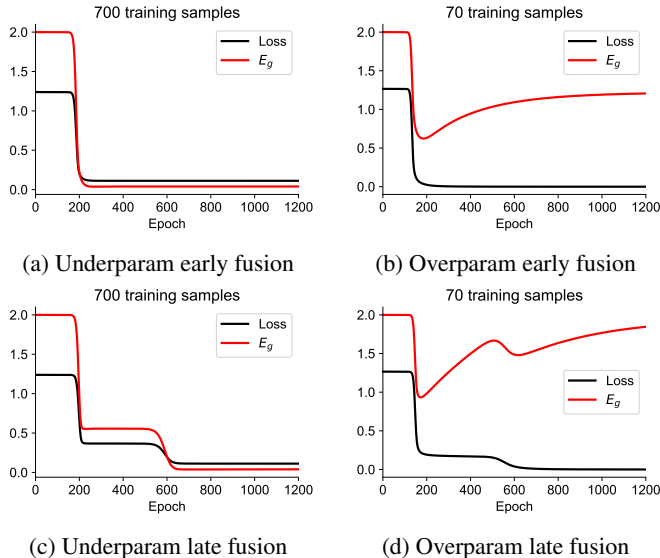


Figure 9: Overparameterized and underparameterized two-layer early and late fusion linear networks. Both modalities have 50 input dimensions with a total of 100. Further implementation details can be found in Appendix I. The input covariance is $\Sigma_A = \mathbf{I}, \Sigma_B = 3\mathbf{I}, \Sigma_{AB} = \mathbf{0}$. The target output is generated as $y = \mathbf{w}^* \mathbf{x} + \epsilon$, where $\mathbf{w}^* = \mathbf{1}/10$ and the noise is independently sampled from $\epsilon \sim \mathcal{N}(0, 0.5^2)$. (a,b) Loss and generalization error trajectories of a two-layer early fusion linear network trained with 700, 70 training samples. (c,d) Loss and generalization error trajectories of a two-layer late fusion linear network trained with 700, 70 training samples.

G.1 UNDERPARAMETERIZED REGIME

Training loss well reflects the corresponding generalization error in the underparameterized regime as shown in Figs. 9a and 9c, where there is sufficient amount of data and thus the training data distribution is close to the true data distribution. The convergent generalization error is lower than the convergent training loss because the training set contains noise. In the underparameterized regime, analysis on the training loss translates to the generalization error.

G.2 OVERPARAMETERIZED REGIME

In the overparameterized regime, the number of samples is insufficient compared to the number of effective parameters, which is the input dimensions for linear networks (Advani et al., 2020).

As shown in Fig. 9b, the overparameterized early fusion linear network learns both modalities during one transitional period. The generalization error decreases during the transitional period and increases afterwards as predicted by theory (Advani et al., 2020). If early stopping is adopted, we obtain a model that has learned from both modalities and not overfitted.

As shown in Fig. 9d, the overparameterized late fusion linear network learns the faster-to-learn modality first and overfits this modality during the unimodal phase when the training loss plateaus but the generalization error increases. In this case, there is a dilemma between overfitting one modality and underfitting the other. If optimal early stopping on the generalization error is adopted, training terminates right after the first modality is learned while the second modality has not, yielding a model with strong and permanent unimodal bias. If training terminates after both modalities are learned, the network has overfitted the first modality, which can yield a generalization error higher than that in early stopping. Thus overfitting is a mechanism that can convert the transient unimodal phase to a generalization deficit and permanent unimodal bias.

As identified with empirical evaluations in Wang et al. (2020), different modalities overfit and generalize at different speeds in multimodal networks. With theoretical evidence, we elaborate on that in late fusion linear network, the difference in speeds gives rise to the unimodal phase in both un-

derparameterized and overparameterized regime and further causes the dilemma between overfitting one modality and underfitting the other in the overparameterized regime.

H TWO-LAYER RELU NETWORKS

H.1 RELU NETWORKS WITH LINEAR TASK

Two-layer ReLU networks are trained to learn the same linear task introduced in Fig. 2. The loss and weights trajectories in Fig. 10 are qualitatively the same as Fig. 2, except that learning evolves about two times slower and the converged weights are two times larger.

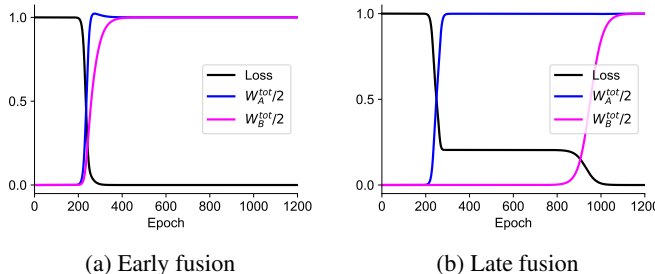


Figure 10: Two-layer early fusion and late fusion ReLU networks trained on a linear tasks. The same dataset, number of neurons, and initialization scale as Fig. 2 are used, except that ReLU activation is added to the hidden layer. (a,b) Time trajectories of loss and total weights in the two-layer early fusion (a) and late fusion (b) ReLU network.

H.2 RELU NETWORKS WITH NONLINEAR TASK

We consider a simple nonlinear task of learning $y = x_A + \text{XOR}(x_B)$, where $x_A \in \mathbb{R}$, $x_B \in \{[1, 1], [1, -1], [-1, 1], [-1, -1]\}$. The term $\text{XOR}(x_B)$ refers to performing exclusive-or to the two dimensions of x_B . We train two-layer early fusion and late fusion ReLU networks to learn this task with variance of the linear modality $\text{Var}(x_A) \equiv \sigma_A = 1, 2$ or 3 . We inspect the features in ReLU networks by plotting the first-layer weights since the features in a two-layer ReLU network lie in its first-layer weights (Xie et al., 2017).

As shown in Figs. 11b, 11d and 11f, two-layer late fusion ReLU networks always solve the task by consistently forming the four perpendicular XOR features. We can see two transitions in the loss trajectories of late fusion ReLU networks, which are similar to the transitions in linear networks. Late fusion ReLU networks learn the XOR modality during one transitional period and learn the linear modality during the other. As an example of converged features shown in Fig. 11d, \mathbf{W}_B^1 has taken on the rank-two XOR structure and \mathbf{W}_A^1 has grown in scale while preserving its independence from \mathbf{W}_B^1 at initialization.

As shown in Figs. 11a, 11c and 11e, two-layer early fusion ReLU networks struggle to extract the XOR features. In the early stage of training, features in \mathbf{W}^1 favor particular directions in the three-dimensional space that are most correlated with the target output in early fusion ReLU networks. In comparison with features in late fusion ReLU networks, the first layer weights for modality A do not preserve its independence from weights modality B as shown in Fig. 11c. The larger the variance of the linear modality, the closer the favorable direction is to the direction of the linear modality. In later stage of training, features in \mathbf{W}^1 can rotate or scatter, giving rise to multiple transitional periods as shown in Fig. 11c and its corresponding video in the supplementary material. For a large variance in the linear modality, the features are highly aligned with the linear modality direction and the network can be stuck in a local minimum, failing to learn from the XOR modality as in Fig. 11e.

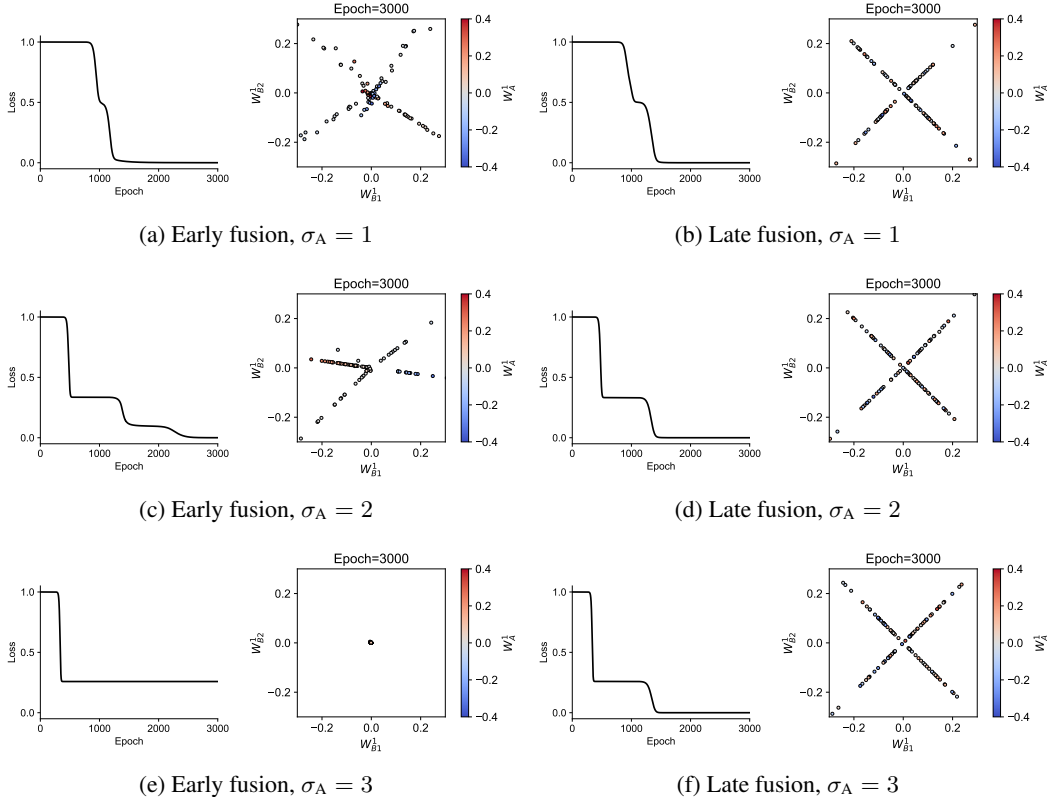


Figure 11: Two-layer early fusion and late fusion ReLU networks trained on a XOR and linear task. The early fusion network has 100 hidden neurons. The late fusion network has 100 hidden neurons in both branches. In every panel, the loss trajectory is plotted on the left and the first-layer weights at the end of training are plotted on the right. The first-layer weights are three-dimensional: the two dimensions of \mathbf{W}_B^1 are plotted as the position of the dots and \mathbf{W}_A^1 is plotted as the color of the dots. (a,c,e) Two-layer early fusion ReLU networks. (b,d,f) Two-layer late fusion ReLU networks. Complementary videos can be found in the supplementary material.

I IMPLEMENTATION DETAILS

I.1 DATA GENERATION

For data generation, we sample input data points from zero mean normal distribution $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ and generate the corresponding target output from a groundtruth linear map $y = \mathbf{w}^* \mathbf{x}$. All datasets do not contain noise except Fig. 9. The number of training samples is 4096 for all experiments except Fig. 9. Hence, all experiments, except Fig. 9, fall into the underparameterized regime where the training loss well reflects the generalization error.

Note that we do not lose generality by using linear datasets because the learning dynamics of linear networks as given in Eqs. (3) and (4) only concern the input correlation matrix Σ and input-output correlation matrix Σ_{yx} . Hence, datasets generated from any distribution and target map with the same Σ, Σ_{yx} will have the same learning dynamics in linear networks.

I.2 MULTIMODAL LINEAR NETWORKS

All early fusion networks we used have 100 neurons in every hidden layer. Late fusion networks have 100 neurons in every hidden layer for both branches. Intermediate fusion networks have 100 neurons in every pre-fusion hidden layer for both branches and 200 neurons in every post-fusion hidden layer. All networks are trained with full batch gradient descent with learning rate $\eta = 0.04$. For Figs. 3b, 3c and 5b to 5d, we run experiments with 5 random seeds and report the average.