

# Critical Learning Periods for Multisensory Integration in Deep Networks

Michael Kleinman<sup>1\*</sup> Alessandro Achille<sup>2</sup> Stefano Soatto<sup>2</sup>  
<sup>1</sup>University of California, Los Angeles <sup>2</sup>AWS AI Labs

michael.kleinman@ucla.edu {aachille, soattos}@amazon.com

## Abstract

We show that the ability of a neural network to integrate information from diverse sources hinges critically on being exposed to properly correlated signals during the early phases of training. Interfering with the learning process during this initial stage can permanently impair the development of a skill, both in artificial and biological systems where the phenomenon is known as a critical learning period. We show that critical periods arise from the complex and unstable early transient dynamics, which are decisive of final performance of the trained system and their learned representations. This evidence challenges the view, engendered by analysis of wide and shallow networks, that early learning dynamics of neural networks are simple, akin to those of a linear model. Indeed, we show that even deep linear networks exhibit critical learning periods for multi-source integration, while shallow networks do not. To better understand how the internal representations change according to disturbances or sensory deficits, we introduce a new measure of source sensitivity, which allows us to track the inhibition and integration of sources during training. Our analysis of inhibition suggests cross-source reconstruction as a natural auxiliary training objective, and indeed we show that architectures trained with cross-sensor reconstruction objectives are remarkably more resilient to critical periods. Our findings suggest that the recent success in self-supervised multi-modal training compared to previous supervised efforts may be in part due to more robust learning dynamics and not solely due to better architectures and/or more data.

## 1. Introduction

Learning generally benefits from exposure to diverse sources of information, including different sensory modalities, views, or features. Multiple sources can be more informative than the sum of their parts. For instance, both views of a random-dot stereogram are needed to extract the

synergistic information, which is absent in each individual view [17]. More generally, multiple sources can help identify latent common factors of variation relevant to the task, and separate them from source-specific nuisance variability, as done in contrastive learning.

Much information fusion work in Deep Learning focuses on the design of the architecture, as different sources may require different architectural biases to be efficiently encoded. We instead focus on the *learning dynamics*, since effective fusion of different sources relies on complex phenomena beginning during the early epochs of training. In fact, even slight interference with the learning process during this *critical period* can permanently damage a network’s ability to harvest synergistic information. Even in animals, which excel at multi-sensor fusion, a temporary deficit in one source during early development can permanently impair the learning process: congenital strabismus in humans can cause permanent loss of stereopsis if not corrected sufficiently early; similarly, visual/auditory misalignment can impair the ability of barn owls to localize prey [18]. In artificial networks, the challenge of integrating different sources has been noted in visual question answering (VQA), where the model often resorts to encoding less rich but more readily accessible textual information [2, 6], ignoring the visual modality, or in audio-visual processing, where acoustic information is often washed out by visual information [32].

Such failures are commonly attributed to the mismatch in learning speed between sources, or their “information asymmetry” for the task. It has also been suggested, based on limiting analysis for wide networks, that the initial dynamics of DNNs are very simple [16], seemingly in contrast with evidence from biology. In this paper, we instead argue that *the early learning dynamics of information fusion in deep networks are both highly complex and brittle, to the point of exhibiting critical learning periods similar to biological systems.*

In Sect. 2, we show that shallow networks do not exhibit critical periods when learning to fuse diverse sources of information, but *deep* networks do. Even though, unlike animals, artificial networks do not age, their learning success is still decided during the early phases of training. The ex-

\*Work conducted during an internship at AWS AI Labs.



Figure 1. **Decomposition of information between different modalities.** Two modalities can have unique information, common information (denoted by the overlap in the venn-diagram), or synergistic information (denoted by the additional ellipse in the right panel). Task-relevant information (shown in red) can be distributed in a variety of ways across the different modalities. Task-relevant information can be mostly present in Modality A (left), shared between modalities (center-left), or could require unique (center-right) or synergistic information from both modalities (right).

istence of critical learning periods for information fusion is not an artifact of annealing the learning rate or other details of the optimizer and the architecture. In fact, we show that critical periods for fusing information are present even in a simple deep linear network. This contradicts the idea that deep networks exhibit trivial early dynamics [16, 23]. We provide an interpretation for critical periods in linear networks in terms of mutual inhibition/reinforcement between sources, manifest through sharp transitions in the learning dynamics, which in turn are related to the intrinsic structure of the underlying data distribution.

In Sect. 3, we introduce a metric called “Relative Source Variance” to quantify the dependence of units in a representation to individual sources, allowing us to better understand inhibition and fusion between sources. Through it, in Sect. 4, we show that temporarily reducing the information in one source, or breaking the correlation between sources, can permanently change the overall amount of information in the learned representation. Moreover, even when downstream performance is not significantly affected, such temporary changes result in units that are highly polarized and process only information from one source or the other. Surprisingly, we found that the final representations in our artificial networks that were exposed to a temporary deficit mirrored single-unit animal representations exposed to analogous deficits (Fig. 4, Fig. 6).

We hypothesize that features inhibit each other because they are competing to solve the task. But if the competitive effect is reduced, such as through an auxiliary cross-source reconstruction task, the different sources can interact synergistically. This supports cross-modal reconstruction as a practical self-supervision criterion. In Sect. 4.4, we show that indeed auxiliary cross-source reconstruction can stabilize the learning dynamics and prevent critical periods. This lends an alternate interpretation for the recent achievements in multi-modal learning as due to the improved stability of the early learning dynamics due to auxiliary cross-modal reconstruction tasks, rather than to the design of the archi-

ture.

Empirically, we show the existence of critical learning periods for multi-source integration using state-of-the-art architectures (Sect. 4.3-4.4). To isolate different factors that may contribute to low-performance on multi-modal tasks (mismatched training dynamics, different informativeness), we focus on tasks where the sources of information are symmetric and homogeneous, in particular stereo and multi-view imagery. Even in this highly controlled setting, we observe the effect of critical periods both in downstream performance and/or in unit polarization. Our analysis suggests that pre-training on one modality, for instance text, and then adding additional pre-trained backbones, for instance visual and acoustic, as advocated in recent trends with Foundation Models, yields representations that fail to encode synergistic information. Instead, training should be performed across modalities at the outset. Our work also suggests that asymptotic analysis is irrelevant for deep network fusion, as their fate is sealed during the initial transient learning. Also, conclusions drawn from wide and shallow networks do not transfer to deep networks in use in practice.

### 1.1. Related Work

**Multi-sensor learning.** There is a large literature on sensor fusion in early development [27], including homogeneous sensors that are spatially dislocated (e.g., two eyes), or time-separated (e.g., motion), and heterogeneous sources (e.g., optical and acoustic, or visual and tactile). Indeed, given *normal learning*, humans and other animals have the remarkable ability to integrate multi-sensory data, such as incoming visual stimuli coming into two eyes, as well as corresponding haptic and audio stimuli. Monkeys have been shown to be adept at combining and leveraging arbitrary sensory feedback information [9].

In deep learning, multi-modal (or *multi-view learning*) learning typically falls into two broad categories: learning a joint representation (fusion of information) and learning an aligned representation (leveraging coordinated informa-

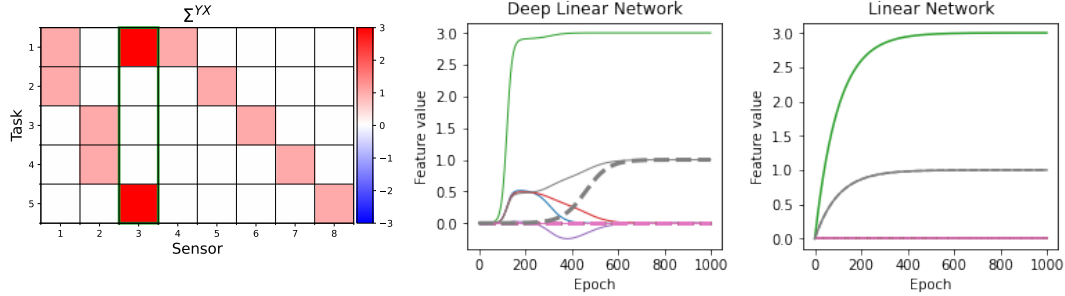


Figure 2. **(Left)**  $\Sigma^{yx}$ , with the highlighted green column representing the sensor that was dropped. **(Center)** We show total weights attributed to each feature (shown in different colors) during training in a deep linear network. The solid lines represent the dynamics when training with all features. The dashed lines represent the behavior when training with the green feature disabled. Note that disabling the green feature prevents the gray feature from being learned during the initial transient **(Right)** Same experiment with a shallow linear network. In this case the learning dynamics of the gray feature perfectly overlap in both cases.

tion in the multiple views) [5]. A fusion-based approach is beneficial if there is synergistic information available in the different views, while an alignment-based approach is helpful if there is shared information common to the different views (Fig. 1). Such a division of information typically affects architectural and model choices: synergistic information requires the information from the different modalities to be fused or combined, whereas shared information often serves as a self-supervised signal that can align information from the different modalities, as in contrastive learning [8, 29, 30], correlation based [3], and information-theoretic approaches [20, 21].

**Critical periods in animals and deep networks:** Such architectural considerations often neglect the impact coming from multisensory learning dynamics, where information can be learned at different speeds from each sensor [34]. Indeed, [33] showed that humans and animals are peculiarly sensitive to changes in the distribution of sensory information early in training, in a phenomenon known as *critical periods*. Critical periods have since been described in many different species and sensory organs. For example, barn owls originally exposed to misaligned auditory and visual information cannot properly localize prey [22]. Somewhat surprisingly, similar critical periods for learning have also been observed in deep networks. [1] found that early periods of training were critical for determining the asymptotic network behavior. Additionally, it was found that the timing of regularization was important for determining asymptotic performance [12], with regularization during the initial stages of training having the most influential effect.

**Masked/de-noising Autoencoders:** Reconstructing an input from a noisy or partial observation has been long used as a form of supervision. Recently, in part due to the successful usage of transformers in language [31] and vision tasks [11], such a pre-training strategy has been successfully applied to text [10] and vision tasks [14]. An extension of this has been recently applied to multi-modal data [4].

**Models of learning dynamics** We consider two approaches to gain analytic insight into the learning dynamics of deep networks. [25, 26] assume that the input-output mapping is done by a deep linear network. We show that under this model critical periods may exist. [16, 23] assume instead infinitely wide networks, resulting in a model linear with respect to the parameters. In this latter case, no critical period is predicted contradicting our empirical observations on finite networks.

## 2. A model for critical periods in sensor-fusion

We want to establish what is the difference, in terms of learning dynamics, between learning how to use two sources of information at the same time, or learning how to solve a task using each modality separately and then merging the results. In particular we consider the counterfactual question: if we disable sensor A during training, would this change how we learn to use sensor B? To start, let's consider the simple case of a linear regression model  $\mathbf{y} = \mathbf{W}\mathbf{x}$  trained with a mean square error loss

$$L = \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \|\mathbf{y}^{(i)} - \mathbf{W}\mathbf{x}^{(i)}\|^2$$

where  $D = \{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$  is a training set of i.i.d. samples. In this simplified setting, we consider each component  $x_k$  of  $\mathbf{x}$  as coming from a different sensor or source. To simplify even further, we assume that the inputs have been whitened, so that the input correlation matrix  $\Sigma^x = \frac{1}{N} \sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)T} = \mathbf{I}$ .

In this case, the learning dynamics of any source is independent from the others. In fact, the gradient of the weight  $w_{jk}$  associated to  $x_k$  and  $y_j$  is given by

$$-\nabla_{w_{jk}} L(\mathbf{W}) = -\nabla_{w_{jk}} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} \|\mathbf{y}^{(i)} - \mathbf{W}\mathbf{x}^{(i)}\|^2 = \Sigma_{jk}^{yx} - w_{jk}$$

and does not depend on any  $w_{hl}$  with  $w_{hl} \neq w_{jk}$ . The answer to the counterfactual question is thus negative in this

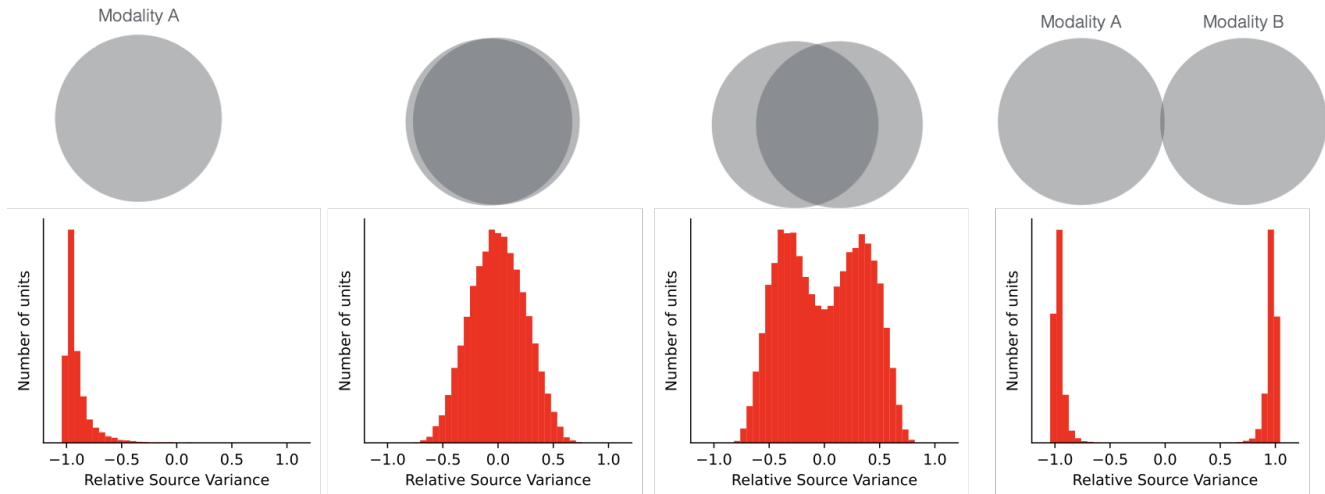


Figure 3. **Example RSV distributions and relation to information diagrams.** (Left) Representations that vary predominantly due to one modality. (Center-Left, Center-right) All units in the representation vary nearly equally with both modalities. (Right) Units in the representation that vary uniquely with each sensor, which is reflected by a polarized RSV distribution.

setting: adding or removing one source of information (or output) will not change how the model learns to extract information from the other sources. However, we now show that the addition of depth, even *without* introducing non-linearities, makes the situation radically different.

To this effect, consider a deep linear network with one hidden layer  $\mathbf{y} = \mathbf{W}^2\mathbf{W}^1\mathbf{x}$ . This network has the same expressive power (and the same global optimum) as the previous model. However, this introduces a mutual dependency between sensors (due to the shared layer) that can ultimately lead to critical periods in cross-sensor learning. To see this, we use an analytical expression of the learning dynamics for two-layer deep networks [25, 26]. Let  $\Sigma^{yx} = \frac{1}{N} \sum_{i=1}^N \mathbf{y}^{(i)} \mathbf{x}^{(i)T}$  be the cross-correlation matrix between the inputs  $\mathbf{x}$  and the target vector  $\mathbf{y}^1$  and let  $\Sigma^{yx} = U S V^T$  be its singular-value decomposition (SVD). [26] shows that the total weight  $\mathbf{W}(t) = \mathbf{W}^2(t)\mathbf{W}^1(t)$  assigned to each source at time  $t$  during the training can be written as

$$\mathbf{W}(t) = \mathbf{W}^2(t)\mathbf{W}^1(t) = \mathbf{U}\mathbf{A}(t)\mathbf{V}^T \quad (1)$$

$$= \sum_{\alpha} a_{\alpha}(t) \mathbf{u}^{\alpha} \mathbf{v}^{\alpha T} \quad (2)$$

where

$$a_{\alpha}(t) = \frac{s_{\alpha} e^{2s_{\alpha}t/\tau}}{e^{2s_{\alpha}t/\tau} - 1 + s_{\alpha}/a_{\alpha}^0}. \quad (3)$$

This leads to non-linear learning dynamics where different features are learned at sharply distinct points in time [26]. Moreover, it leads to entanglement between the learning dynamics of different sources due to the eigenvectors  $\mathbf{v}^{\alpha}$  mixing multiple sources.

<sup>1</sup>Note that  $\mathbf{W} = \Sigma^{yx}$  is also the global minimum of the MSE loss  $L = \frac{1}{N} \sum_i \frac{1}{2} \|\mathbf{y}^{(i)} - \mathbf{W}\mathbf{x}^{(i)}\|^2$ .

Disabling (or adding) a source of information corresponds to removing (or adding) a column to the matrix  $\Sigma^{yx}$ , which in turn affects its singular-value decomposition and the corresponding learning dynamics. To see how this change may affect the learning dynamics, in Fig. 2 we compare the weights associated to each sensor during training for one particular task. In solid we show the dynamics with all sensors active at the same time. In dashed line we show the dynamics when one of the sensor is disabled. We see that disabling a sensor (green in the figure) can completely inhibit learning of other task-relevant features (e.g., the gray feature) during the initial transient. This should be compared with the learning dynamics of a shallow one-layer network (Fig. 2, right) where all task-relevant features are learned at the same time, and where removal of a source does not affect the others.

In deep linear networks, the suboptimal configuration learned during the initial transient is eventually discarded, and the network reverts to the globally optimal solution. In the following we show this is not the case for standard non-linear deep networks. While the initial non-trivial interaction between sources of information remain, the non-linear networks are unable to unlearn the suboptimal configurations learned at the beginning (owing to the highly non-convex landscape). This can result in permanent impairments if a source of information is removed during the initial transient of learning, which reflects the trends observed in critical periods in animals.

### 3. Single Neuron Sensitivity Analysis

Before studying the empirical behavior of real networks on multi-sensor tasks, we should consider how to quantify the effect of a deficit on a down-stream task. One way

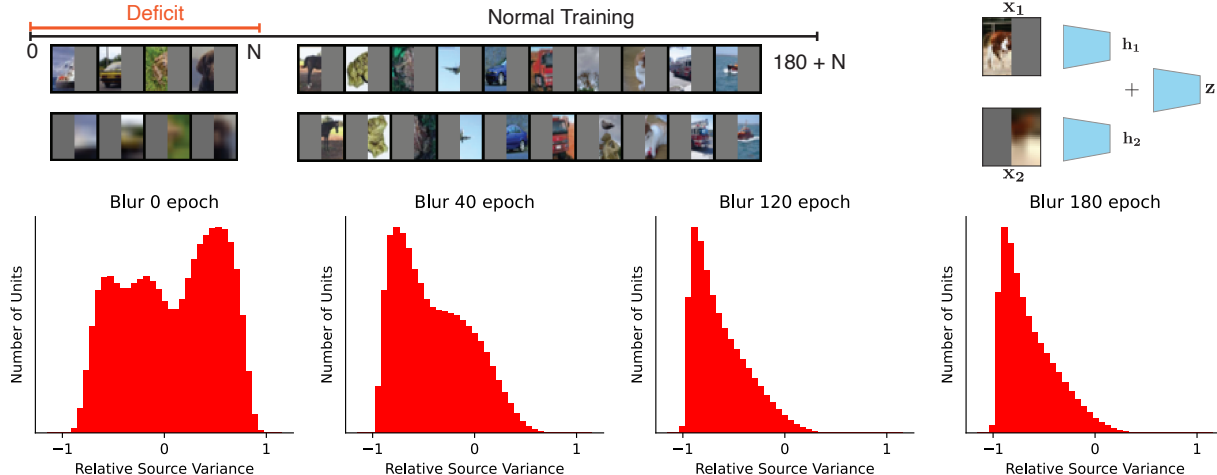


Figure 4. **Experimental setup and sensor selectivity as a function of a blurring deficit length.** (Top) In our experiments, we train the network with a deficit (blurred images to one pathway shown here) for the first  $N$  epochs, and then continue training with normal images for 180 more epochs. We feed each half of an image to the early stages of a ResNet-18, and then additively combine the representations from both pathways (followed by stages of common processing). We refer to this architecture as *Split-ResNet*. (Bottom) RSV distribution of units in last layer representation  $z$  for increasing duration of deficit (blur to one pathway) after resumption of normal training. With a sufficiently long deficit, the units in the representation remain only sensitive to the initially uncorrupted pathway, and do not vary with the initially corrupted pathway.

is to look at the final performance of the model on the task. For example, animals reared with a monocular deprivation deficit have reduced accuracy on a visual acuity test and, similarly, deep networks may show reduced classification accuracy [1]. However, in some cases deficits may not drastically impair the accuracy but may still affect how the model is organized internally. Individuals with strabismus or ambliopia can perform just as well on most tasks, since the individual information coming from each sensor separately is enough to compensate. But the connectivity scheme of the synapses may change so that neurons eventually process only information from one sensor or the other, and not from both together, as observed in individuals without deficits [33].

To understand whether units in a representation of multisensory inputs depend on both sensors or only a particular sensor, we introduce a measure of *Relative Source Variance*. We first define the *Source Variance* ( $SV$ ) for unit  $i$  of a representation due to sensor  $A$ , conditioned on an example  $b$  as

$$SV_i(A, b) = \text{Var}(f(A, B)_i | B = b), \quad (4)$$

where  $f$  denotes the mapping from multisensory inputs to the representation and  $i$  indexes the unit of the representation. We note that the value of  $SV_i(A, b)$  depends on the example  $b$ . We use an analogous formula for  $SV_i(B, a)$ .

Typically, we are interested in the distribution of the Source Variance of the units  $i$  in a representation, as a function of many examples  $a$  and  $b$ . To capture this, we define a

notion of Relative Source Variance (RSV) for unit  $i$  as:

$$RSV_i(a, b) = \frac{SV_i(A, b) - SV_i(B, a)}{SV_i(A, b) + SV_i(B, a)} \quad (5)$$

If the RSV is 1, this means that the unit is only sensitive to sensor  $A$ , and if the RSV is  $-1$ , the unit is sensitive to sensor  $B$ . To compute  $SV(A, b)$  (and analogously for  $SV(B, a)$ ) from samples, we fix a sample  $b$ , and vary the inputs  $a$ , sampling from  $a \sim p(a)$ . We run this for multiple fixed samples from  $b$ , performing the computation over a batch. We perform analogous computations for  $SV(B, a)$ . We compute the  $RSV_i(a, b)$  for all units  $i$  from a representation, and for many examples  $a$  and  $b$ . We then plot the distribution of RSVs, aggregating across all units (see, e.g., Fig. 4-6). In particular, we track how the distribution changes as a result of sensory deficits and perturbations, as well as how the distribution changes during normal training. Note that  $-1 \leq RSV_i(a, b) \leq 1$ . If  $RSV_i(a, b) = 1$  (or  $-1$ ) is 1, this means that the unit is only sensitive to sensor  $A$  (or  $B$ ). If  $RSV_i(a, b) = 0$  the unit is equally sensitive to both sensors. For controlled simulations (See Appendix A.1), we show the variety of distributions of units in a representation that the RSV can measure in Fig. 3. In our experiments, described next, we computed the RSV on units from the final layer before linear classification.

#### 4. Critical learning periods in deep multi-sensor networks

In this section, we investigate the learning dynamics of deep networks during the initial learning transient when

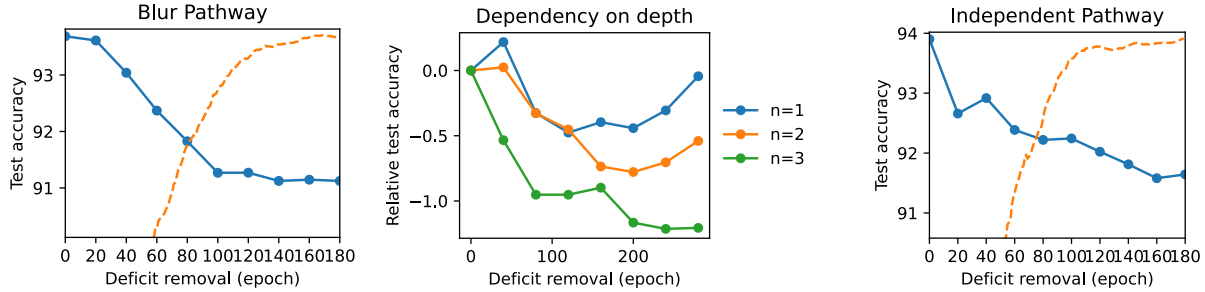


Figure 5. **Decrease in downstream performance as a function of the deficit length.** (Left) Final test accuracy (blue) when applying a blurring deficit to one pathway of Split-ResNet. Even though the network is exposed to a subsequent number of uncorrupted paired observations, the network cannot later learn to optimally fuse the information. The orange dashed line represents accuracy of a normal network during training. (Center) The effect of a deficit is most pronounced when increasing the depth of the network (see Appendix for architecture detail). (Right) We also observe a degradation of test performance using a dissociation deficit (feeding uncorrelated views). We note that the effect is less marked than the blurring, due to better ability to compensate. Additional runs are shown in Fig. 16.

multiple source of information are present. We evaluate how temporary perturbations of the relation between the two sensors during the training can change the final outcome. To exclude possible confounding factors, in all our experiments, the two input sources are perfectly symmetrical (same data distribution and same informativeness for the task) which ensures that any asymmetry observed in the final model is due to the perturbation.

#### 4.1. Inhibition of a weak source

Uncorrected vision problems in an eye during early childhood can cause permanent visual impairment in humans, whereas even after correction the patient only sees through the unaffected eye and does not recover vision in the affected eye (amblyopia, or lazy-eye). We explore whether such inhibition of a sensor can happen in DNNs following a similar experimental setup to [1]. To simulate binocular data from single images, we partition each image in a left and right crop and feed each to two separate pathways of the network, which are then fused in an additive manner at a later stage. For each initial pathway, we used the early stages of a ResNet-18 backbone. We then simulate the blurry vision of a weak eye by downsampling the input of the right pathway by  $4\times$ , and then resized the image to the original size. After training for  $t_0$  initial epochs with the blur deficit, we remove it and train for further 180 epochs to ensure convergence (see Appendix for details). Our experimental setup is schematized in Fig 4 (top panel). Here we focus on the simple CIFAR-10 classification dataset, and we later examine different architectures and datasets, and learning approaches.

At the end of the training, both sensors are working well and contain partially disjoint information about the task variable, so the network would benefit from using both of them. However, in Fig. 4 we see from the RSV that weakening the right sensor by blurring it during the initial transient will permanently inhibit its use even after re-

moving the deficit. More specifically, at the end of normal training units in the network attend equally to either sensor (leftmost panel). However, in the network trained with a short deficit the neurons only encode information about the “initially good” left sensor (the RSV of the units concentrates around -1, rightmost panel). This mirrors the ocular dominance findings present in monkeys with a cataract [33, Fig. 7]. Similarly, the longer the deficit is present during the initial training, the more the downstream performance on the CIFAR-10 classification task is impaired (Fig. 5, left). However, the reduction of performance is not as drastic as the RSV change, since the network can compensate and achieve a good accuracy on the task using only the good sensor.

**Dependency on depth.** In Sect. 2 we note that depth is fundamental to make critical periods emerge in multi-sensor networks. We further claim that increasing the depth of the network makes critical periods more evident. Indeed, in Fig. 5 (center) we show that increasingly deeper network have increasingly more marked permanent impairment as a result of a temporary deficit.

#### 4.2. Learning synergistic information

We have seen that temporary weakening of one sensor may completely inhibit its learning. We now consider an alternative deficit where the two sensors are both working well, but are initially trained on uncorrelated data and only later trained together. This situation is common in every day machine learning, for example when pre-training backbones on different modalities separately (e.g., a text and a vision backbone) and then fine-tuning them together on a downstream task.

**Dissociation deficit.** To keep the two modalities symmetrical, we consider a similar set up as before where we feed to each pathway of a network the left and a right crop of an image. Both crops are now always full-resolution. However, we introduce a *dissociation* deficit, during which

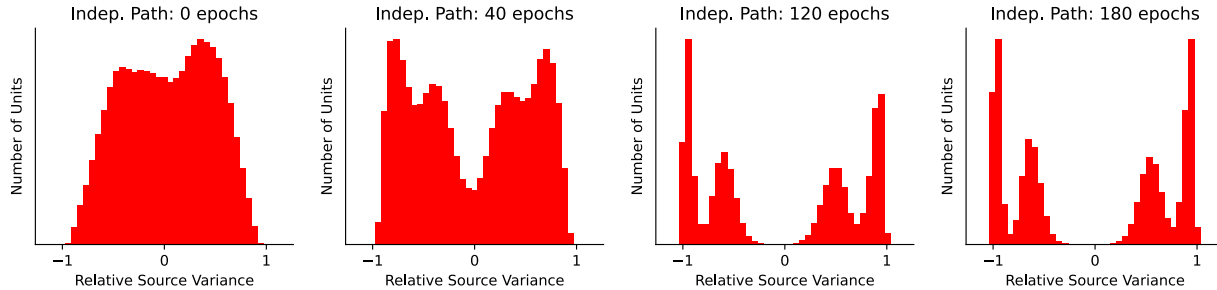


Figure 6. **Sensor selectivity as a function of a dissociation deficit length.** We examined the asymptotic representations and found that, when exposed to a sufficiently long deficit of broken correlations between the views, the network could no longer learn a bimodal distribution that learned common features, but instead resulted in a polarized representation in which units are sensitive to either view (but none to both).

the right crop is sampled from a different image than the left one. During the dissociation, the task is to predict either the class of the left image or the right image with probability 0.5. This deficit removes any synergistic information between the two pathways, but still encourages the two pathways to extract any unique information from the inputs.

We observe that this setup too has a critical period: In Fig. 6, we see that, after normal training, the units are equally sensitive to both the left and right inputs (histogram clusters around zero). However, after training with an increasingly longer dissociation deficit, the histogram becomes increasingly polarized around  $\pm 1$ , suggesting that each unit is encoding information only about the right or the left image. This precludes the possibility that the network is extracting synergistic information from the two views (which would entail units that process information from both sensors). This mirrors the ocular dominance representations observed in strabismic monkeys [33, Fig. 10-12]. Similarly to the dissociation deficit, in strabismus, the eyes are not aligned, thus breaking the normal correlation between the views. The dissociation deficit also produces a permanent impairment in the downstream performance (Fig. 5, right) but again the effect is not as drastic as in the RSV plot since the network compensates by using each pathway separately (albeit synergistic information is lost).

### 4.3. Synergistic information in videos

So far we have seen that supervised deep networks, similar to humans and animals, have critical periods for learning correspondences between multi-view data. We confirmed this both at the behavioural (measured in terms of performance and visual acuity for the deep networks and animals respectively) and at the representation level, quantified by the neuron sensitivity. We now investigate whether such phenomenon generalize across learning strategies, architecture, and tasks.

**Multi-View Transformer.** Aside from integrating information from different sensors, animals and artificial networks need to be able to integrate information through time.

We can think of frames of a video as being different views or sources of information that are correlated through time, and we can study how a network learns to integrate such information. We opted to use a more flexible transformer-based visual architecture, which has recently achieved state-of-the-art results in computer vision tasks [11, 14], and language tasks [10, 31]. Visual transformers are typically trained either with a supervised loss [11] or a masking-based objective, followed by fine-tuning [14]. We focus now on the first case, and analyze the second in the next section. In order to process multiple frames of a video, we use a modified Multi-Modal Masked Auto-Encoder [4], which we train in a fully supervised fashion. We refer to this as a *Multi-View Transformer*.

To capture multiple views of a scene, we opted to use the the Kinetics Action classification video dataset [7], which consist in classifying one of 400 possible actions given a video clip. To adapt the task to our setting, from each video we select two random frames that are a multiple of 0.33 seconds apart to comprise our two views, and feed them to the *Multi-View Transformer*. Due to their temporal correlation, the two frames together contain more information (the motion) than either frame individually. We use a similar dissociation deficit as in the previous section: During the dissociation deficit period, we sampled the two frames from independent videos in order to break their temporal correlation. In this case, the classification label coming from either view with  $p = 0.5$  (see Appendix for training details). We introduce the deficit in a sliding window of fixed size starting, and vary the starting time to measure the sensitivity of each part of the training process.

Even on a largely different architectures (transformer instead of ResNet) and a more complex task (action classification on natural video instead of CIFAR-10), in Fig. 8 we observe the same trends as in the previous section. Training with a temporary dissociation deficit permanently prevents the network from extracting synergistic temporal information from the frames. Unlike in the previous experiment, since the synergistic information is fundamental for the ac-

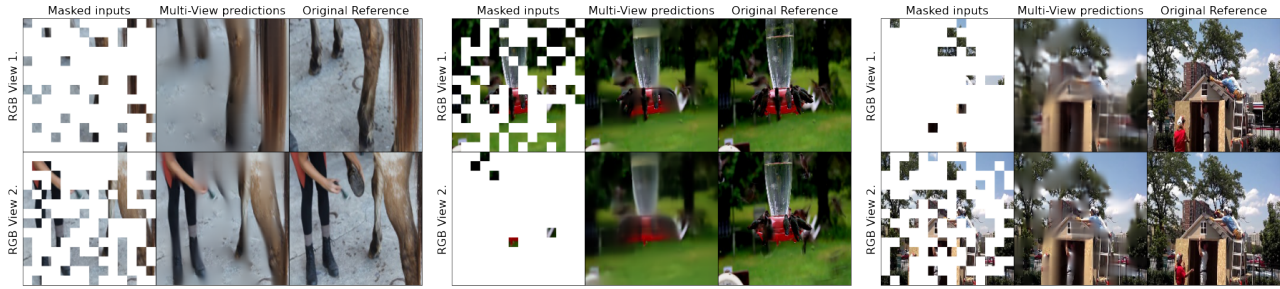


Figure 7. **Top.** Example inputs (left column), reconstructions (middle columns), and original targets (right columns) for the Multi-View Transformer, with random sampling of patches from the two views. Note that the model can reconstruct missing information from one view using the other.

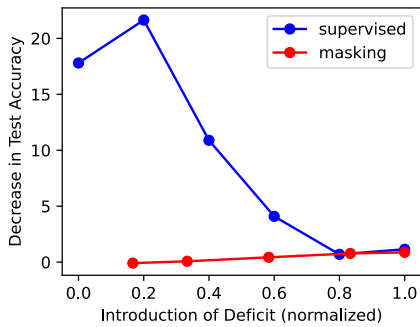


Figure 8. **Masking objective with cross-sensor reconstruction loss does not exhibit a critical learning period.** We found that the unsupervised network was much more robust to perturbations early in the training (red trace), whereas that supervised objective was not (blue trace).

tion classification task, the network cannot compensate the deficit and perturbations during the critical period also results in an harsh decrease of up to 20% in the final test accuracy (Fig. 8, blue trace).

#### 4.4. Overcoming critical periods with cross-sensor reconstruction

Our previous experiments suggest that critical periods can be caused by competition between sensors which increases the selectivity of the units. If this is the case, we may hypothesize that training adding a cross-sensor reconstruction objective may help forcing the unit to learn how to encode cross-sensor information. To test this hypothesis, we train the Multi-View transformer of Sec. 4.3 using the cross-sensor masking-based reconstruction objective of [4] and compare it with the supervised case. The self-supervised masked-image reconstruction task could encourage correspondences to be learned (if un-occluded parts of one view are helpful for reconstructing the other view), and may force learning synergistic information irrespective of the initial transient. In Fig. 7, we show that indeed the masking-based pre-training is successful in using information from one source to predict masked patches of the other.

We train using the same protocol as Sec. 4.3 to pre-train the Multi-View Transformer using the cross-reconstruction objective. We then subsequently fine-tuned for 20 epochs on the downstream supervised classification task (see Appendix for details). In Fig. 8 we see that the unsupervised network was much more robust to perturbations early in the training, whereas that supervised objective was not. To understand whether such robustness was due to large changes to the representation when fine-tuning, we applied the RSV on the output of the encoder’s representation and found that while the resulting distribution became slightly more symmetrically balanced, it retained a similar bimodal distribution to the pre-trained representation. (Fig. 14).

## 5. Discussion

We have shown – in a variety of architectures and tasks – the existence of critical learning periods for multi-source integration: a temporary initial perturbation of an input source may permanently inhibit that source, or prevent the model from learning how to combine multiple sources. These trends replicate similar phenomena in animals, and point to the underlying complexity and brittleness of the learning dynamics that allow a network (or an animal) to fuse information. To simplify the analysis of the learning dynamics, we focused on tasks with homogeneous sources (stereo, video). We leave to future work to further study the role played by the asymmetry between sources (e.g., different informativeness or ease). Our theoretical and empirical analysis leads to several suggestions: Pre-training different backbones separately on each modality, as advocated in some foundational models, may yield representations that ultimately fail to encode synergistic information. Instead, training should be performed across modalities at the outset. On the theoretical side, our work suggests that analysis “at convergence” of the learning dynamics of a network are irrelevant for sensor fusion, as their fate is sealed during the initial transient learning. It also suggests that conclusions drawn from wide and shallow networks may not transfer to deep networks in current use.



## References

- [1] Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep networks. In *International Conference on Learning Representations*, 2019. 3, 5, 6, 12
- [2] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. Analyzing the behavior of visual question answering models. *arXiv preprint arXiv:1606.07356*, 2016. 1
- [3] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *International conference on machine learning*, pages 1247–1255. PMLR, 2013. 3
- [4] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multimaie: Multi-modal multi-task masked autoencoders. *arXiv preprint arXiv:2204.01678*, 2022. 3, 7, 8, 13
- [5] 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, 2018. 3
- [6] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32, 2019. 1
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 7, 13
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3
- [9] Maria C Dadarlat, Joseph E O’doherly, and Philip N Sabes. A learning-based approach to artificial sensory feedback leads to optimal integration. *Nature neuroscience*, 18(1):138–144, 2015. 2
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3, 7
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 7
- [12] Aditya Sharad Golatkar, Alessandro Achille, and Stefano Soatto. Time matters in regularizing deep networks: Weight decay and data augmentation affect early learning dynamics, matter little near convergence. In *Advances in Neural Information Processing Systems 32*, pages 10677–10687. Curran Associates, Inc., 2019. 3
- [13] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 13
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021. 3, 7
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 12
- [16] Wei Hu, Lechao Xiao, Ben Adlam, and Jeffrey Pennington. The surprising simplicity of the early-time learning dynamics of neural networks. *Advances in Neural Information Processing Systems*, 33:17116–17128, 2020. 1, 2, 3
- [17] Bela Julesz. Binocular depth perception of computer-generated patterns. *Bell System Technical Journal*, 39(5):1125–1162, 1960. 1
- [18] Eric R Kandel, James H Schwartz, Thomas M Jessell, Steven A Siegelbaum, and A J Hudspeth. *Principles of neural science*. McGraw-Hill, New York, fifth edition, 2013. 1
- [19] Michael Kleinman, Alessandro Achille, Daksh Idnani, and Jonathan Kao. Usable information and evolution of optimal representations during training. In *International Conference on Learning Representations*, 2021. 13, 14
- [20] Michael Kleinman, Alessandro Achille, Stefano Soatto, and Jonathan Kao. Gacs-korner common information variational autoencoder. *arXiv preprint arXiv:2205.12239*, 2022. 3
- [21] Michael Kleinman, Alessandro Achille, Stefano Soatto, and Jonathan C Kao. Redundant information neural estimation. *Entropy*, 23(7):922, 2021. 3
- [22] Eric I Knudsen and Phyllis F Knudsen. Sensitive and critical periods for visual calibration of sound localization by barn owls. *Journal of Neuroscience*, 10(1):222–232, 1990. 3
- [23] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019. 2, 3
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 13
- [25] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013. 3, 4
- [26] 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. 3, 4, 12
- [27] Linda Smith and Michael Gasser. The development of embodied cognition: Six lessons from babies. *Artificial life*, 11(1-2):13–29, 2005. 2
- [28] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014. 12
- [29] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European conference on computer vision*, pages 776–794. Springer, 2020. 3

- [30] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020. [3](#)
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [3](#), [7](#)
- [32] Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal classification networks hard? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12695–12705, 2020. [1](#)
- [33] Torsten N Wiesel. Postnatal development of the visual cortex and the influence of environment. *Nature*, 299(5884):583–591, 1982. [3](#), [5](#), [6](#), [7](#)
- [34] 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 *International Conference on Machine Learning*, pages 24043–24055. PMLR, 2022. [3](#)
- [35] Yilun Xu, Shengjia Zhao, Jiaming Song, Russell Stewart, and Stefano Ermon. A theory of usable information under computational constraints. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. [13](#)