

# BLOCK-RECURRENT DYNAMICS IN ViTs

**Mozes Jacobs\***  
Harvard University

**Thomas Fel\***  
Harvard University

**Richard Hakim\***  
Harvard University

**Alessandra Brondetta**  
Osnabrück University

**Demba Ba**  
Harvard University

**T. Andy Keller**  
Harvard University

## ABSTRACT

As Vision Transformers (ViTs) become standard backbones across vision, a mechanistic account of their computational phenomenology is now essential. Despite architectural cues that hint at dynamical structure, no settled framework interprets Transformer depth as a well-characterized flow. In this work, we introduce the **Block-Recurrent Hypothesis (BRH)**, arguing that trained ViTs admit a block-recurrent depth structure such that the computation of the original  $L$  blocks can be accurately rewritten using only  $k \ll L$  distinct blocks applied recurrently. Across diverse ViTs, between-layer representational similarity matrices suggest few contiguous phases. To determine whether these phases reflect reusable computation, we operationalize our hypothesis in the form of block recurrent surrogates of pretrained ViTs, which we call **Recurrent Approximations to Phase-structured TransfORMers (Raptor)**. Using small-scale ViTs, we demonstrate that phase-structure metrics correlate with our ability to accurately fit Raptor and identify the role of training and stochastic depth in promoting the recurrent block structure. We then provide an empirical existence proof for BRH in foundation models by showing that we can train a Raptor model to recover 96% of DINOv2 ImageNet-1k linear probe accuracy in only 2 blocks while maintaining equivalent runtime. To provide a mechanistic account of these observations, we leverage our hypothesis to develop a program of **Dynamical Interpretability**. We find *(i)* directional convergence into class-dependent angular basins with self-correcting trajectories under small perturbations, *(ii)* token-specific dynamics, where cls executes sharp late reorientations while patch tokens exhibit strong late-stage coherence reminiscent of a mean-field effect and converge rapidly toward their mean direction, and *(iii)* a collapse of the update to low rank in late depth, consistent with convergence to low-dimensional attractors. Altogether, we find that a compact recurrent program emerges along the depth of ViTs, pointing to a low-complexity normative solution that enables these models to be studied through principled dynamical systems analysis.

## 1 INTRODUCTION

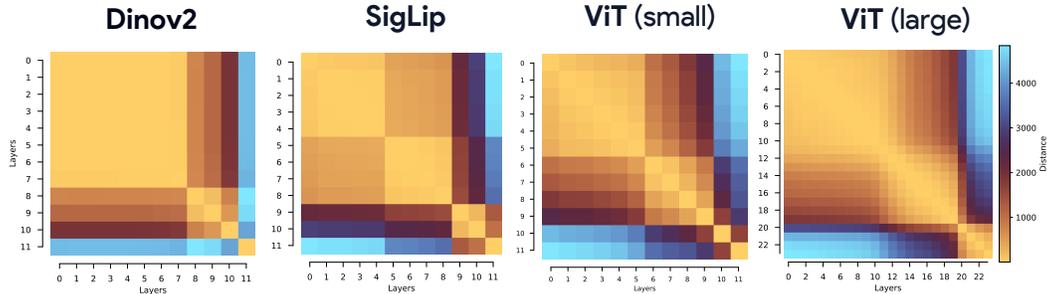
In the last decade, Transformers have become the default neural network architecture across machine learning communities, scaling favorably with data and compute (Vaswani et al., 2017; Kaplan et al., 2020). In particular, Vision Transformers (ViTs) (Dosovitskiy et al., 2020) have become the core architecture used in visual foundation modeling frameworks such as DINOv2 (Oquab et al., 2023; Darcet et al., 2023) and CLIP (Radford et al., 2021); and have come to dominate a wide range of visual tasks, from general visual feature extraction (He et al., 2021; Chiu et al., 2024; Yun, 2025), to diffusion (Peebles & Xie, 2023), image segmentation (Kirillov et al., 2023; Liu et al., 2024), and video processing (Arnab et al., 2021; Baldassarre et al., 2025). This increasing breadth of use motivates a move from empirical optimization to principled understanding.

Two pressures make this understanding urgent. First, safety-critical deployments (Wang & Chung, 2022; Alecu et al., 2022) demand mechanisms whose internal computation is inspectable (Losch et al., 2021), diagnosable (Adebayo et al., 2020), and verifiable (Tjeng & Tedrake, 2019) rather than

\* Equal contribution.

Correspondence to {mozesjacobs, tfel, rhakim, takeller}@g.harvard.edu.

 [kemptoner.github.io/raptor](https://github.com/kemptoner/kemptoner.github.io/raptor)



**Figure 1: Layer-layer similarity matrices across diverse Vision Transformers reveal block-structure.** Despite differences in scale and training objectives, all models exhibit contiguous block structure along depth, visible as phase-segmented regions of high similarity. Beyond representational similarity, this raises the question of whether a deeper *functional* recurrence underlies these patterns, hinting at block-wise reusability of computation across layers. In this work, we investigate this hypothesis, showing that these phase segments correspond to blocks with functional similarity, which can be approximated by a single shared block applied recurrently along depth.

opaque. As these models proliferate across domains, the ability to explain (Doshi-Velez & Kim, 2017; Gilpin et al., 2018; Kim et al., 2018), manipulate, and verify their behavior becomes increasingly essential. Second, from a scientific inference perspective (Cichy & Kaiser, 2019), understanding what makes these models work is essential for explaining their success. Their strong performance suggests they have discovered effective strategies, and identifying these strategies could advance our broader understanding of visual intelligence. Independently of any comparison to human vision, the goal is to uncover the principles that make these systems so effective.

One promising path toward such understanding is to search for underlying simplicity. Multiple approaches explore different facets of this simplicity, whether in functional expressivity (Montúfar et al., 2014; Telgarsky, 2015; Serra et al., 2018; Balestrierio et al., 2018), symmetry (Cohen & Welling, 2014; Olah et al., 2020), or computation (Wilson, 2025; Goldblum et al., 2023; Schmidhuber, 1997; Mingard et al., 2025). Discovering such simplicity principle should improve both development (Bronstein et al., 2017; Frankle & Carbin, 2019) and interpretability (Bereska & Gavves, 2024; Carvalho et al., 2019; Fel, 2024; Ghorbani et al., 2017; Fel et al., 2023; Smilkov et al., 2017; Sundararajan et al., 2017; Zeiler & Fergus, 2014; Templeton et al., 2024; Bricken et al., 2023). Depth offers a concrete place to look for this simplicity. Residual connections have long suggested a link to dynamical systems (Liao & Poggio, 2020; Veit et al., 2016; Greff et al., 2016; Boulch, 2017; Haber & Ruthotto, 2017), hinting at implicit recurrence even when layers have distinct parameters. This convergent evidence makes plausible a form of algorithmic parsimony (Ma et al., 2022) in which a small set of blocks may be reused across many layers, trading parameters for iterations (Schwarzschild et al., 2021). Related perspectives support this view (Dingle et al., 2020). More concretely, residual updates invite a discrete-time interpretation of depth (Chen et al., 2018; Chalvidal et al., 2020; Sander et al., 2022), attention induces coupled token dynamics (Lu et al., 2019; Geshkovski et al., 2023), and in language models contiguous block recurrence has been observed and exploited (Geiping et al., 2025; Fernando & Guitchounts, 2025; Dehghani et al., 2018; Tan et al., 2023). This renewed interest in recurrent computation (Jolicoeur-Martineau, 2025; Venkatraman et al., 2025) and algorithmic complexity (Shaw et al., 2025; Dingle et al., 2020) as frameworks for understanding neural network simplicity bias further motivates our investigation. However, no existing framework characterizes depth in ViTs as representational flow or determines whether apparent phases correspond to functional reuse. Furthermore, vision explainability research (Bach et al., 2015; Fong & Vedaldi, 2017; Novello et al., 2022; Muzellec et al., 2024; Petsiuk et al., 2018; Hedström et al., 2022; Fel et al., 2025; Gorton, 2024; Kowal et al., 2024; Bau et al., 2017; Vilas et al., 2023) has not leveraged dynamical systems analysis to model emergent network structure.

In this work, we consider recurrence as a candidate source of simplicity and advance the Block-Recurrent Hypothesis (BRH): after training, the depth of a ViT organizes into a few contiguous phases such that the computation of the original  $L$  layers can be rewritten by reusing only  $k \ll L$  distinct blocks applied recurrently. Our starting point is an empirical observation: layer-layer representational similarity matrices consistently exhibit block-diagonal structure across disparate models. However, representational similarity does not necessarily translate to functional equivalence; therefore, we ask: *Does this phase structure reflect genuinely reusable computation?*

**Our contributions.** Our study proceeds in three parts:

- **Empirical evidence for block-recurrent structure.** We demonstrate across diverse ViTs that layer-layer representational similarity matrices exhibit distinct contiguous phases of computation, formalized through the Block-Recurrent Hypothesis. We develop a max-cut algorithm to systematically identify phase boundaries and show that this block structure is *(i)* associated with compressibility by a recurrent architecture and *(ii)* strengthened by stochastic depth.
- **Constructive verification via recurrent surrogates.** We operationalize the BRH by training weight-tied block-recurrent approximations of pretrained ViTs, termed Raptor. Critically, our goal is not compression or efficiency optimization per se, but rather to demonstrate that functional reuse is genuinely possible. Raptor reconstructs the complete internal representation trajectory across all layers, not merely the final output, providing strong evidence for true computational equivalence rather than input-output mimicry. Specifically, we provide empirical evidence for the BRH on foundational vision models by training a Raptor that recovers 96% of DINOv2’s ImageNet-1k linear-probe accuracy using only 2 recurrent blocks, and 98% with 3 blocks.
- **Dynamical systems analysis framework.** Leveraging our hypothesis, we propose a program of Dynamical Interpretability that treats ViT depth as the discrete-time unfolding of an underlying dynamical system, such that the evolution of representations across layers can be analyzed through its temporal dynamics. Our analysis reveals: *(i)* directional convergence into class-dependent angular basins with self-correcting trajectories under perturbations, *(ii)* token-specific dynamics where cls tokens execute sharp late reorientations while patch tokens exhibit strong coherence reminiscent of mean-field behavior, and *(iii)* collapse of layer-to-layer updates to low-rank subspaces consistent with convergence to low-dimensional attractors.

As a first step, we characterize emergent phases in representation space, motivating the BRH.

## 2 EMERGENT PHASE STRUCTURE & THE BLOCK-RECURRENT HYPOTHESIS

Our investigation starts with a simple experiment: we construct layer-layer similarity matrices by computing the cosine similarity of each token at layer  $l$  with the same token at layer  $m$ . As Figure 1 shows, despite varying tasks, objectives, and scales, all models exhibit consistent block-wise organization: high mutual similarity within contiguous blocks and lower similarity across boundaries. This finding is not new and echoes early (and more recent) observations in residual networks (Kornblith et al., 2019; Raghu et al., 2021; Nguyen et al., 2022; Hoang et al., 2025), but raises a fundamental question: does representational similarity reflect deeper computational structure? In fact, representational similarity alone provides no guarantee of functional equivalence. Layers might produce similar representations through entirely different computational pathways, or conversely, functionally equivalent computations might yield representations that appear dissimilar due to linear transformations or noise. The critical question is whether these apparent phases reflect genuine functional recurrence – reusing the same computational operations across layers within each phase. We formalize this possibility through the Block-Recurrent Hypothesis:

**Definition 1** (Block-Recurrent Hypothesis (BRH)). *We consider  $\mathbf{f}$  to be a trained Vision Transformer with nominal depth  $L$  and intermediate maps  $\mathbf{f}_\ell : \mathcal{X} \rightarrow \mathcal{A}_\ell$ ,  $\ell \in \{1, \dots, L\}$ . We say that  $\mathbf{f}$  satisfies the  $\varepsilon$ -BRH if for any  $\ell$ , there exist  $k \ll \ell$  blocks  $\mathbf{B}_1, \dots, \mathbf{B}_k$  and integers  $n_1, \dots, n_k$  with  $\sum_{j=1}^k n_j = \ell$  such that:*

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} (\|\mathbf{f}_\ell(\mathbf{x}) - (\mathbf{B}_k^{(n_k)} \circ \dots \circ \mathbf{B}_1^{(n_1)})(\mathbf{x})\|_F) \leq \varepsilon,$$

where  $\mathbf{B}_j^{(n_j)}$  denotes  $n_j$  repeated applications of the same parameter-tied block  $\mathbf{B}_j$  and the entire approximation maintaining equivalent computational cost.

Here,  $\|\cdot\|_F$  denotes the Frobenius norm and  $\mathbb{P}$  is a probability distribution over natural images. To put it simply, this hypothesis states that a ViT’s  $L$  layers can be replaced by  $k \ll L$  recurrent blocks that reproduce the entire internal trajectory at equivalent computational cost. By requiring intermediate layer fidelity rather than just final output matching, we rule out trivial solutions where computation is concentrated in a single block. The constraint  $k \ll L$  with parameter tying ensures genuine functional reuse rather than simple parameter copying.

To test this hypothesis, we first operationalize it by proposing a method to construct these approximations using recurrent architectures, developing a specialized training technique described below.

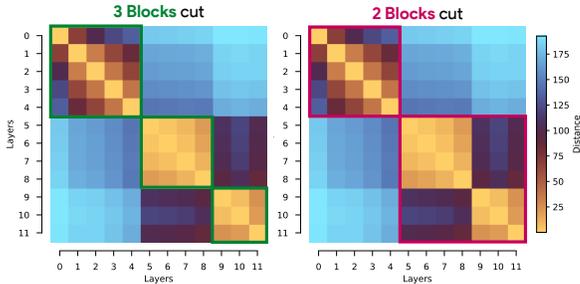


Figure 2: **Block discovery via max-cut segmentation of the layer–layer similarity matrix.** Our algorithm partitions depth into contiguous segments by maximizing within-block similarity and minimizing cross-block cosine similarity. Shown are two cuts of the same ViT-B: with 3-blocks (left, green) and 2-blocks (right, magenta). These cuts reveal candidate block boundaries where the representation dynamics undergo sharp transitions, providing an operational method for detecting contiguous recurrent phases in trained ViTs.

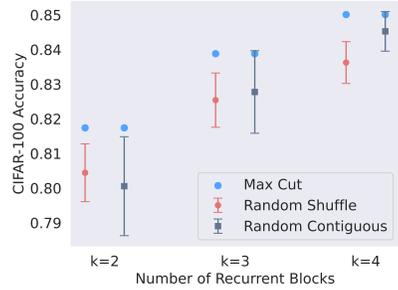


Figure 3: **Evaluation of Raptor models on CIFAR-100 using our max-cut partitioning algorithm versus random partitions.** Reported values are classification accuracy. Results for random partitions are aggregated over 10 different random partitions. Random shuffle refers to non-contiguous (fragmented) random partitions.

**Operationalizing Block-Recurrence with Raptor.** Since the BRH asserts only the existence of recurrent blocks satisfying its conditions and not their precise form, the most direct validation is constructive: demonstrating existence by example of a recurrent model approximating  $f$ .

However, recurrent architectures are notoriously difficult to train: rollouts can drift as small state errors compound over steps, and gradients propagated through many recurrent applications often vanish or explode (Pascanu et al., 2013; Linsley et al., 2020; De et al., 2024). Standard backpropagation through time becomes unstable as recurrence depth increases, and the model must learn to simultaneously handle both the forward dynamics and its own prediction errors in closed loop (Williams & Zipser, 1989). To circumvent these challenges, we will leverage the intermediate layer activations as training targets, enabling a staged approach that combines the stability of teacher forcing with the self-consistency required for autoregressive deployment. More precisely, we introduce a procedure to distill existing Vision Transformers into Recurrent Approximations to Phase-structured TransfORMers (Raptors), using  $k$  parameter-tied blocks with repetition counts determined by a max-cut phase discovery algorithm. This approach transforms the abstract hypothesis into a concrete architectural and training framework that can be empirically validated.

This constructive approach requires that Raptor models reproduce the internal activations of the full ViT they approximate, similar to Dasgupta & Cohn (2025); Sanh et al. (2019); Shleifer & Rush (2020), not merely mimic the final output<sup>1</sup>. The BRH implies that such reproduction should be possible within tolerance  $\varepsilon$ , making activation matching a natural training objective. Formally, let  $f$  be a reference ViT with intermediate activations  $\mathbf{a}_\ell(\mathbf{x}) \equiv \mathbf{f}_\ell(\mathbf{x}) \in \mathbb{R}^{t \times d}$  for  $\ell = 0, \dots, L$ , where layer  $\ell = 0$  denotes the patch encoder and  $1 \leq \ell \leq L$  refer to transformer layers. Here,  $t$  is the number of tokens and  $d$  the feature dimension. Let  $\mathbf{B}_j$  denote the  $j$ -th parameter-tied block in our recurrent decomposition. The Raptor approximation produces activations:

$$\tilde{\mathbf{a}}_\ell(\mathbf{x}) \equiv (\mathbf{B}_k^{(n_k)} \circ \dots \circ \mathbf{B}_1^{(n_1)})(\mathbf{a}_0(\mathbf{x})) \tag{1}$$

where the composition covers layers 1 to  $\ell$  according to our phase segmentation. We train Raptor using an autoregressive loss (AR) that enforces trajectory fidelity across all intermediate layers:

$$\mathcal{L}_h^{\text{AR}}(\mathbf{x}) = \mathbb{E}_{\mathbf{x}} \left( \sum_{\ell=1}^h \|\tilde{\mathbf{a}}_\ell(\mathbf{x}) - \mathbf{a}_\ell(\mathbf{x})\|_F \right), \quad h \leq L. \tag{2}$$

This approach teaches each block to approximate its designated segment, while the overall model reproduces the teacher networks’ complete representational trajectory. However, this formulation does not specify how to determine the block boundaries or phase assignments. We address this now with a simpler algorithmic approach based on the representational similarity structure observed earlier.

<sup>1</sup>Unlike classical distillation, which typically supervises logits (and occasionally a few intermediate “hints”), we enforce one-to-one alignment of all layers representations across the entire depth for the same inputs. The recurrent surrogate must generate the teacher’s intermediate activations, not just its predictions.

**Choosing partitions.** For a given number of blocks  $k$ , we must introduce a practical method to determine the number of recurrent iterations of each block ( $n_k$ ); in other words, where the recurrent “phases” of computation begin and end. We accomplish this by casting this “block discovery” process as a weighted max-cut problem [Goemans & Williamson \(1995\)](#), solved via dynamic programming (see [Subsection A.1](#) for details). Specifically, the algorithm seeks to partition depth into contiguous segments by maximizing within-block similarity and minimizing cross-block similarity. We visualize the results of this procedure applied to ViT-B in [Figure 2](#), demonstrating that the discovered blocks align reasonably with qualitative assessment.

To validate this approach, we train recurrent transformer models using max-cut partitions to reproduce the activations of trained Vision Transformers on CIFAR-100 (validation set accuracy 90.7%). Remarkably, as shown in [Figure 3](#), Raptor student models with only 2 recurrent blocks closely match the performance of the ViT-B teacher models they approximate. The max-cut algorithm provides partitions that achieve strong performance, with accuracy often greater than one standard deviation above randomly chosen partitions. This observation suggests that the representational block-similarity structure is closely associated with functional block-recurrent phases.

To assess how unique each block-recurrent layer is and control for degenerate solutions to functional block-recurrence, we tested whether swapping a layer from a different block could substitute for a layer within a target block ([Figure 15](#)). Using DINOv2-B trained on ImageNet-1k, we find that while intra-block swapping preserves accuracy, inter-block swapping leads to model collapse. These results demonstrate that representational block-similarity structure predicts functional block-recurrent structure, and that layer identity is functionally unique to each block.

**How do blocks emerge?** Having operationalized the BRH and demonstrated a method for block discovery, we now turn to the mechanistic origins of this phenomenon: *under what conditions does this block-recurrent structure emerge in trained Vision Transformers?* To investigate this systematically, we examine small-scale ViTs where we control training conditions and isolate potential contributing factors. Specifically, we hypothesize that training and stochastic depth ([Huang et al., 2016](#)) may promote the emergence of block-recurrent patterns.

Motivated by evidence that residual networks tolerate variable effective depth ([Wu et al., 2019](#)), we examined the effect of stochastic depth (SD) on block recurrence. During training, each layer is dropped independently with probability  $p$ , applied uniformly across depth. We trained ViT-B/14 from random initialization on CIFAR-100, using the cls token for the linear probe across a sweep of SD  $p$  rates. We observe an increase in layer-layer similarity with increasing SD  $p$  rates ([Figure 4A, E](#)). Next, we used these trained ViT networks as teachers for student Raptor models (see [Appendix A](#)). Raptor models were trained to reconstruct the hidden activation states of the ViT teacher across layers. Raptor forward passes are fully autoregressive, meaning each layer’s output is fed into the next layer and is also trained to match the corresponding layer in the teacher network ([Figure 10](#)). We quantify the similarity of the cls and patch token representations in each layer between the teacher and student networks as the  $R^2$  of their matched token embeddings ([Figure 4B](#)).

We observe that, as ViT stochastic depth increases, a separately trained Raptor student model becomes significantly better at reconstructing the ViT’s layerwise hidden states ([Figure 4D](#)). In addition, we observe that increasing SD improves CIFAR-100 classification accuracy in both ViT teacher and Raptor student networks ([Figure 4C](#)). We combine the above results in [Figure 4E](#) and observe a strong positive association between the ViT’s layer-layer representational similarity and Raptor reconstruction fidelity. These results demonstrate that stochastic depth regularization increases layer-layer similarity and recurrent compressibility.

We also sought to determine whether recurrent compressibility was dependent on several other factors, including whether the ViT teacher was trained or untrained, and whether skip branches are included in the network architecture. We observe that network hidden layer activations in untrained ViTs can be reconstructed by Raptor better than trained ViTs. Interestingly, we observe that removing the skip branch from an untrained ViT only slightly reduces reconstruction accuracy. Finally, we trained Raptor students on checkpoints from a single ViT allowed to overfit (via reduced weight decay) on CIFAR-100 classification ([Figure 13](#)). We observe that Raptor reconstruction accuracy ( $R^2$ ) remains high until the ViT network begins to overfit ([Figure 13B-D](#)), after which cls token reconstruction and classification accuracy drops. These results add to the results on stochastic depth and demonstrate how the natural architecture of ViTs as well as the normative state of properly regularized ViT networks demonstrate emergent block recurrent dynamics.

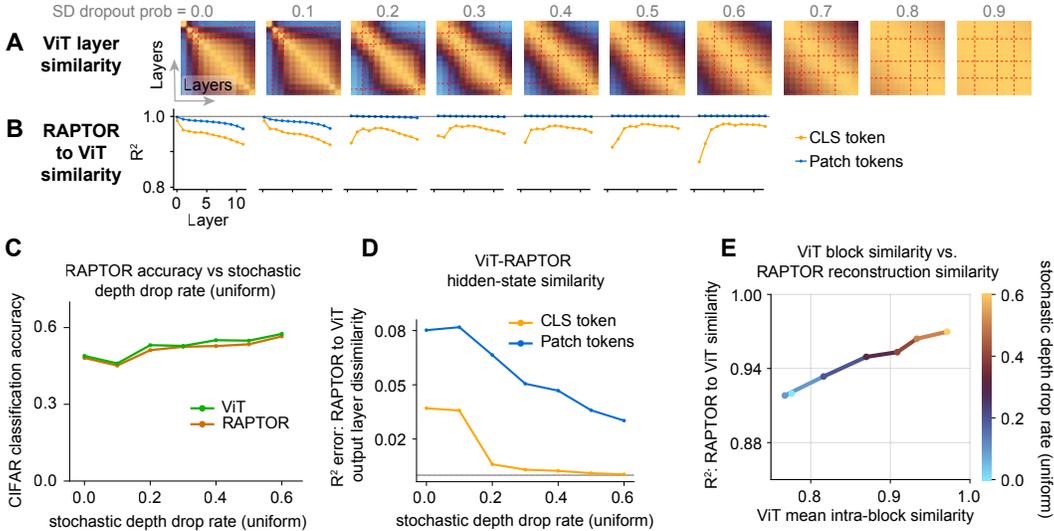


Figure 4: **Stochastic depth promotes representational similarity across layers block-recurrence.** **A)** ViT layer-layer cosine similarity matrices for models trained with increasing stochastic depth (SD) dropout probability  $p$  (probabilities of 0.0-0.9, uniform over layer depth). Dashed red lines delineate blocks, as defined by the max-cut algorithm. Higher SD  $p$  values lead to a more similar representation across layers. **B)** Layerwise teacher-student representational alignment  $R^2$  (Raptor vs. ViT) of the class CLS and patch tokens. Increases in SD  $p$  correspond to an increase in the ability of Raptor to match the ViT’s layerwise representations. ViT models for SD=0.7-0.9 show aberrant training dynamics and are excluded from this and further analysis. **C)** CIFAR classification accuracy for a ViT trained on CIFAR, and a Raptor model with  $k = 3$  blocks trained to match the hidden state of the ViT. **D)** Last layer hidden-state similarity  $R^2$  error (equivalent to  $1 - R^2$ ) of the ViT and Raptor model as a function of SD  $p$ . Increases in stochastic depth lead to a greater ability to reconstruct ViT function using Raptor. **E)** Association between layer-layer representational similarity and Raptor reconstruction  $R^2$ . Stochastic depth encourages the formation of more similar blocks of layers within the ViT, which facilitates approximation by the recurrent Raptor model.

Taken together, these toy-model experiments support the view that the observed representational block structure reflects both an intrinsic property of the residual ViT architecture and a learned/emergent functional recurrence. Using the methods established here, we next scale up our application of Raptor to modern large-scale foundation models.

### 3 SCALING Raptor TO FOUNDATION MODELS

Having demonstrated the BRH on controlled experiments, we now test whether it extends to large-scale foundation models. We apply Raptor to DINOv2, chosen for its widespread adoption across vision tasks, and optimize it to reproduce DINOv2’s internal activations on ImageNet-1k.

**Architecture and Training.** For all experiments, we use ImageNet-1K and extract activations from a pretrained DINOv2 (ViT-Base) model, applying our max-cut algorithm to identify  $k \in \{2, 3, 4\}$  recurrent block partitions. Each recurrent block  $B(\cdot)$  mirrors the block of DINOv2 (for detail about the block see Section A). We train Raptor using a two-stage approach that combines teacher forcing (TF) and autoregressive (AR) objectives. In the first stage, teacher forcing trains each block to predict the immediate next layer given the correct previous layer, while the autoregressive objective requires the model to use its own predictions as inputs for subsequent layers. Refer to Figure 11 for this hybrid training algorithm. This approach yields the following total loss:

$$\mathcal{L}_{\text{total}}(\mathbf{x}) = \lambda \mathcal{L}_{\text{TF}}(\mathbf{x}) + (1 - \lambda) \mathcal{L}_{\text{AR}, H}(\mathbf{x}) + \Omega(\boldsymbol{\theta}),$$

where  $\Omega(\boldsymbol{\theta})$  denotes additional regularization applied to each tied block. See appendix A for complete details. An attentive reader will notice that this training approach naturally lends itself to parallelization: since each block operates on a distinct layer range, the first training stage can be executed simultaneously across multiple GPUs or machines. Each block learns to approximate

its designated segment of the original network using the combined objective, with teacher forcing gradually annealed to zero as training progresses. The first stage thus allows blocks to develop their specific computational roles while benefiting from ground-truth activations as inputs.

The second stage connects all trained blocks into the complete recurrent architecture and trains the entire system end-to-end using only the autoregressive loss (i.e.  $\lambda = 0$ ). This crucial phase teaches blocks to coordinate their computations and handle their own predicted activations rather than relying on ground-truth inputs from the teacher network. The transition from teacher forcing to pure autoregression ensures that the final model can operate independently while maintaining fidelity to the original network’s representational trajectory.

We provide an implementation framework at <https://kempnerinstitute.github.io/raptor>. With the training methodology established, we now evaluate how effectively Raptors can reproduce the performance of their teacher networks across multiple vision tasks.

**Results.** We evaluate Raptor against DINOv2 by training linear probes on ImageNet-1k (classification), ADE20k (semantic segmentation), and NYUv2 (monocular depth), covering both classification and dense prediction. For ImageNet-1k, we initialize the classifier from the public DINOv2 probe and report the best score across initialization and fine-tuning. In all experiments, the ViT backbone is frozen for both Raptor and DINOv2; only the linear heads are updated, and we reuse DINOv2’s patch embedding and final layer normalization (both also frozen). Results appear in Table 2. Raptor performs well across tasks and is stronger on classification: with  $k = 3$  it attains 83.0% top-1 on ImageNet-1k (about 98% of DINOv2 ViT-B and above ViT-S; see Fig. 5), with little deviation across runs ( $\sigma = 0.1$ ). Accuracy improves markedly from  $k = 2$  to  $k = 3$  and then saturates at  $k = 4$ . In short, a two-block Raptor retains about 96% of DINOv2 ViT-B with a frozen backbone, a compact rewriting that substantiates the BRH. This conclusion is further supported by Figure 14, which shows that Raptor maintains a high cosine similarity to DINOv2 activations through depth, confirming that it captures the representational dynamics of the original model.

**Ablations.** Although our aim is not maximal compression nor exact accuracy matching, we perform targeted ablations to identify the factors most critical to Raptor performance (Table 1). Training with teacher forcing alone (Stage 1 only), while computationally efficient, leads to complete collapse with poor accuracy ( $\sim 3\%$  on ImageNet-1K), indicating that one-step supervision is insufficient without exposure to the full autoregressive trajectory. Introducing the autoregressive loss and gradually annealing teacher forcing to zero raises accuracy by more than 68%, underscoring the necessity of closed-loop training for stable block-recurrent approximation. Further gains come from depth scaling (Appendix A.3), where each block has a learned vector embedding of its target layer index, making Raptor a non-autonomous dynamical system (i.e., the update rule explicitly depends on the iteration count rather than state alone). Up-weighting the final block’s c1s token loss yields further improvements (Eq. 3, Appendix A). Finally, connecting all blocks and fine-tuning the model end-to-end with the autoregressive objective (i.e. the second stage) dramatically increases performance, and a final boost is obtained by fine-tuning the linear probe. Now that we have shown that the BRH holds for a foundation model, and before turning

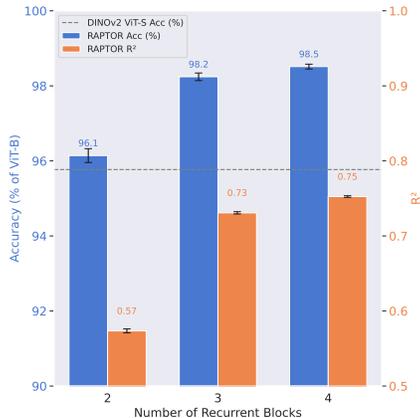


Figure 5: **Raptor’s performance on ImageNet-1k as a function of DINOv2 ViT-B accuracy (left), and  $R^2$  score (right).** DINOv2 ViT-S accuracy shown as a dashed horizontal line. Results are aggregated over three model runs trained on different random seeds.

Method	Accuracy
Teacher Forcing (TF)	3.9
+ Autoreg (anneal TF)	72.7 $\uparrow$ 68.8
+ Depth Scaling	75.2 $\uparrow$ 2.5
+ Weighted c1s	76.7 $\uparrow$ 1.5
+ Second Stage	82.4 $\uparrow$ 5.7
+ Finetune (Classifier)	83.0 $\uparrow$ 0.6

Table 1: **Ablations to original Raptor(k=3) model,** showing ImageNet-1k accuracy with DINOv2 pretrained linear classifier. Second Stage refers to putting all three blocks together and training the full model autoregressively.

Up-weighting the final block’s c1s token loss yields further improvements (Eq. 3, Appendix A). Finally, connecting all blocks and fine-tuning the model end-to-end with the autoregressive objective (i.e. the second stage) dramatically increases performance, and a final boost is obtained by fine-tuning the linear probe. Now that we have shown that the BRH holds for a foundation model, and before turning

Method	Arch.	IN-1k (Acc $\uparrow$ )	ADE20k (mIoU $\uparrow$ )	NYUv2 (RMSE $\downarrow$ )
Raptor	$k = 2$	$81.2 \pm 0.2$	$39.6 \pm 0.6$	$0.648 \pm 0.003$
	$k = 3$	$83.0 \pm 0.1$	$43.0 \pm 0.3$	$0.618 \pm 0.006$
	$k = 4$	$83.2 \pm 0.1$	$43.6 \pm 0.1$	$0.607 \pm 0.006$
DINOv2	ViT-S	80.9	44.6	0.600
	ViT-B	84.5	47.5	0.578

Table 2: **Performance of Raptor compared to DINOv2 with linear probes.** We report top-1 accuracy on ImageNet-1k, mean Intersection-over-Union (mIoU) on ADE20k semantic segmentation, and root mean squared error (RMSE) on NYUv2 depth estimation. Higher values are better for accuracy and mIoU, while lower values are better for RMSE. Results for Raptor are aggregated over three model runs, each trained with a different random seed, and displayed as  $\mu \pm \sigma$ . For Raptor, *Arch* denotes the number of recurrent blocks, while for DINOv2, *Arch* denotes the ViT backbone.

to Dynamical Interpretability, we first examine one implication of this phenomena: the algorithmic and computational implications of the BRH.

**Algorithmic and computational implications.** At scale, BRH holds in practice: a two-block Raptor recovers most of DINOv2 ViT-B, with three blocks essentially closing the gap. This reveals a strong simplicity bias in trained ViTs: depth reuses a small set of computations, effectively trading parameters for iterations. This reuse has two immediate consequences. First, it compresses the program description length that realizes the network’s computation, suggesting low algorithmic complexity. However, the implication is more subtle than standard Kolmogorov-complexity (Kolmogorov, 1965). While Kolmogorov compression can replace a long program with an arbitrarily short one that runs in unbounded time, Raptor crucially preserves computational cost: applying the same block  $n_j$  times achieves equivalent runtime to  $n_j$  distinct untied blocks. In other words, ViTs admit a more compact program representation *under the same runtime budget*, aligning more closely with Levin’s complexity  $K_{\text{Levin}}$  (Levin, 1973). At scale, BRH holds in practice: a two-block Raptor recovers most of DINOv2 ViT-B, and three blocks close the gap. This reveals a simplicity bias in trained ViTs: depth reuses computations, trading parameters for iterations.

**Claim 1** (BRH guarantees low Levin complexity). *Let  $\mathbf{f}_\ell$  satisfy 0-BRH with  $k \ll \ell$  tied blocks  $\{\mathbf{B}_j\}_{j=1}^k$  and schedule  $(n_j)_{j=1}^k$  with  $\sum_j n_j = \ell$ . Let  $R(\cdot)$  denote block runtime and define the untied teacher runtime  $R(\mathbf{f}_\ell) := \sum_{j=1}^k n_j R(\mathbf{B}_j)$ ; assume runtime parity  $R(\mathbf{B}_j) \leq (1 + \delta)R(\mathbf{B}_\mathbf{f})$ . Then*

$$K_{\text{Levin}}(\mathbf{f}_\ell) \leq \sum_{j=1}^k DL(\theta(\mathbf{B}_j)) + O(k \log \ell) + \log R(\mathbf{f}_\ell) + O(1).$$

See Appendix E for details. Theoretically, BRH guarantee low Levin complexity (compact algorithmic descriptions at unchanged computational cost). Our empirical validation of BRH on DINOv2 then suggest that foundation vision models are algorithmically simpler than their nominal architecture suggests. This compression reinforces emerging evidence that simplicity principles govern successful neural networks (Goldblum et al., 2023; Valle-Perez et al., 2019; Huh et al., 2023). For interpretability research, this offers an encouraging perspective: there exist representational lenses under which seemingly complex models reveal underlying simplicity. High-performing ViTs discover and iteratively reuse a compact set of algorithmic primitives, which is a structural regularity that may provide tractable entry points for mechanistic understanding.

We now pursue one such entry point. Since ViTs compress to recurrent blocks applied iteratively, we propose to analyze their computation as discrete-time dynamical systems, in the next section, we develop this *Dynamical Interpretability* framework to extract insights from the recurrent structure.

#### 4 FROM BLOCK RECURRENCE TO *Dynamical Interpretability* IN ViTs

Having observed block-structured representational similarity and confirmed that this similarity translates to functional recurrence, even in foundational models, we are inclined to seriously consider ViTs as dynamical systems that can be interpreted using dynamical systems analysis tools (Schmid, 2022; Ostrow et al., 2023; Huang et al., 2025) – what we term *dynamical interpretability*. We begin by establishing the basic dynamical properties of this depth flow, and present three key findings: (i) tokens converge directionally toward angular attractors with self-correcting dynamics, (ii) different token types exhibit specialized dynamics with punctuated transitions at phase boundaries, and (iii) later layers exhibit low-rank collective motion under weak contraction, reminiscent of mean-field processes with collapsing update dimensionality.

**Directional Convergence and Angular Attractor Geometry.** We begin by isolating direction and scale. Feature norms increase steadily across depth, precluding Euclidean notions of convergence or attractors; we therefore normalize states and study their angular evolution (Fig. 12). Concretely, let

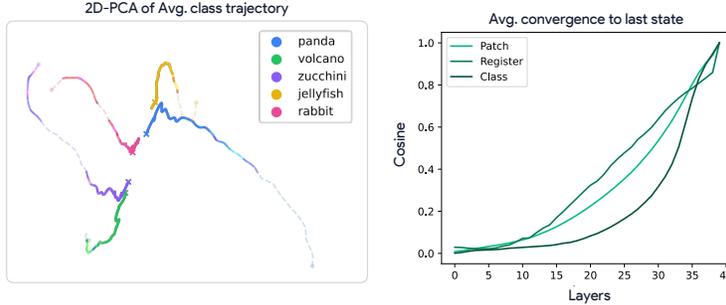


Figure 6: **Directional convergence on the unit sphere.** (Left) Average normalized trajectories (PCA space) collapse into compact class-dependent basins, consistent with low-dimensional angular attractors. (Right) Cosine similarities to final representations  $\gamma_\ell$  are S-shaped, saturating near 1 for cls, registers, and patch, indicating directional fixed points.

$\hat{x}_\ell = \mathbf{x}_\ell / \|\mathbf{x}_\ell\|$  denote the direction of a token at layer  $\ell$ , and consider the depth trajectory  $\{\hat{x}_\ell\}_{\ell=0}^L$  on the unit sphere  $\mathbb{S}^{d-1}$ . Directional convergence is quantified by  $\gamma_\ell = \langle \hat{x}_\ell, \hat{x}_L \rangle$ . Empirically,  $\gamma_\ell$  follows smooth S-shaped curves that approach 1 and saturate in late layers for all token types (Fig. 6, right). This behavior indicates a directional fixed point: while norms grow, directions stabilize so that  $\hat{x}_{\ell+1} \approx \hat{x}_\ell$  as  $\ell$  increases. The late acceleration of  $\gamma_\ell$  suggests phase-local attraction strengthening in the final phase, clarified by our coherence study (Fig. 9, middle). A complementary geometric view comes from projecting the depth trajectories onto a low-dimensional subspace. PCA reveals that sample-specific paths enter class-dependent basins in a shared angular subspace; across 1,000 images from 5 imagenet classes, trajectories curl into compact terminal regions rather than scattering (Fig. 6, left). We interpret these regions as angular attractors: small sets on  $\mathbb{S}^{d-1}$  toward which iterates of the phase-local map steer directions, up to within-class variability. Finally, we probe stability by injecting a small additive perturbation at layer  $\ell$  and following the perturbed direction thereafter. The average perturbed path bends back toward the unperturbed trajectory, indicating local self-correction and on-sphere contraction around the limiting direction (Fig. 8). Together, these measurements establish property (i): token directions evolve under depth toward angular attractors with mild contraction, making directional geometry an appropriate lens for subsequent dynamical analysis.

**Token-Specific Dynamics.** Token groups follow distinct angular update laws. For a token with normalized state  $\hat{x}_\ell$ , define the per-layer angular speed  $s_\ell = \arccos\langle \hat{x}_{\ell+1}, \hat{x}_\ell \rangle$ . Aggregating  $s_\ell$  by token type reveals stable small speeds for registers, intermediate speeds for patches, and sharp late reorientations for cls (Fig. 7). The variance of  $s_\ell$  is smallest for registers after early depth; conversely, cls shows increased late angular activity, reflecting its global aggregator function. These token-specific laws vary across depth. Angular speed statistics display abrupt changes aligned with discovered block boundaries, creating a punctuated pattern where phases maintain near-stationary behavior that is reset at phase transitions (Fig. 7). This structure matches the block-recurrent view in which a phase applies a reused update map with stable statistics before handing off to a new regime at the boundary. Sensitivity analyses corroborate this specialization. Inject a small additive perturbation of magnitude  $\varepsilon$  at layer  $\ell$  and measure the final angular deviation using the cosine distance  $d_{\cos}(\hat{x}_L^{(\varepsilon, \ell)}, \hat{x}_L)$ . The scaled sensitivity  $|\varepsilon|^{-1} d_{\cos}$  decays approximately log-linearly with deeper injection for patch tokens, indicating on-sphere attenuation within phases, whereas it increases for cls when injected late, indicating accumulation at the read-out stage where global information is consolidated (Fig. 8B). Directional convergence rates mirror these roles. Tracking  $c_\ell = \langle \hat{x}_\ell, \hat{x}_L \rangle$  reveals registers approach their terminal directions earliest, patches follow with a smoother rise, and cls saturates only in the final phase where its reorientation peaks (Fig. 6, left). Together, these measurements demonstrate ViT depth implements specialized, phase-local dynamics with phase transitions, consistent with block-recurrent computation.

**Low-Rank Collective Motion and Linearized Depth Flow.** We analyze layer-to-layer token-wise angular updates and observe a progressive collapse to low dimensions. For token-wise angular updates (see App. C). Stable rank and effective rank decrease with depth to near six in the final phase, indicating subspace confinement (Fig. 9, left). Concurrently, patch-token coherence  $\kappa_\ell$  rises sharply

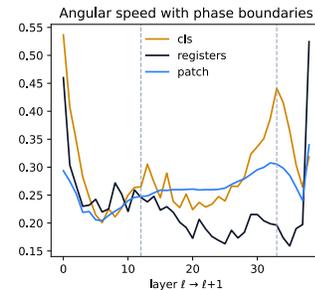


Figure 7: **Token-specific angular speed with phase overlays.** Mean angular speed  $s_\ell$  across depth for cls, registers, and patch, with max-cut phase boundaries from Sec. 2 overlaid as vertical lines.

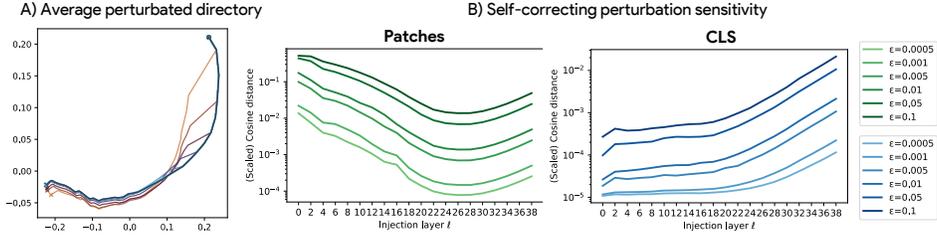


Figure 8: **Self-correction under small angular perturbations.** Perturbed trajectories bend back toward the baseline path, evidencing local basin stability. Sensitivity decays approximately logarithmically with remaining depth for patch tokens, but grows late for cls, consistent with stronger late-stage aggregation.

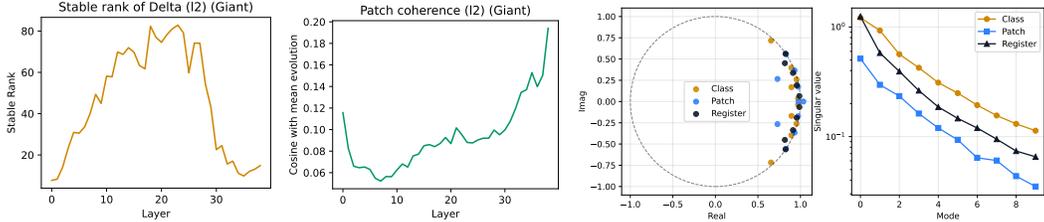


Figure 9: **(Left) Low-rank updates and coordinated patch motion.** Left Stable and effective rank of the layer-to-layer update matrix collapse with depth, indicating confinement to a restricted subspace. Right Patch-token coherence with their mean update direction rises strongly, revealing increasing collective alignment. **(Right) Dynamic Mode Decomposition (DMD) of depth dynamics.** For each token group (cls, registers, patch), we average token states within the group and fit the exact-DMD (see Section D). Each layer state is  $\ell_2$ -normalized to unit norm (trajectories on the unit sphere), so eigenvalue angles  $\arg(\lambda_i)$  characterize angular updates, while radii  $|\lambda_i|$  measure contraction on the sphere (not absolute feature-norm growth). The DMD eigenvalues  $\{\lambda_i\}$  lie just inside the unit circle (dashed) and concentrate near the positive real axis, indicating near-neutral, mostly angular updates with mild on-sphere contraction. cls modes lie closest to +1 (longest memory), registers are slightly more dispersed, and patch shows the widest angular spread and stronger contraction. The cls spectrum decays slowest (highest effective rank/complexity), registers are intermediate, and patch decays fastest (lower-rank dynamics). Together, these spectra support a weakly contracting, block-recurrent depth flow with token-specific complexity.

and peaks late, showing increasingly aligned, collective updates (Fig. 9, middle). The joint pattern (rank collapse with rising coherence) marks a transition from many weakly independent directions to a few shared directions. We then linearize the depth flow via exact DMD on group-averaged,  $\ell_2$ -normalized states, yielding  $\bar{x}_{\ell+1} \approx \bar{A} \bar{x}_\ell$  with rank  $r = 10$  (App. D). Eigenvalues are concentrated just inside the unit circle and near the positive real axis, consistent with weak on-sphere contraction and predominantly angular updates; cls modes lie closest to +1 (longest memory), registers are intermediate, and patches show wider angular spread and stronger contraction (Fig. 9, right). Stacked-depth singular spectra mirror this ordering, decaying slowest for cls and fastest for patches. These results indicate that late depth implements low-rank, near-neutral dynamics that compress variation into a small set of collective directions while preserving long-memory channels for cls.

## 5 DISCUSSION

We advanced the Block-Recurrent Hypothesis (BRH), showing empirically and constructively (via weight-tied surrogates) that recurrence can match untied baselines, and we developed *Dynamical Interpretability* by viewing depth as a flow on directions. This revealed (i) directional convergence to angular attractors with self-correction, (ii) token-specific, phase-local dynamics with punctuated transitions, and (iii) a late low-rank regime that coordinates updates to low dimensional subspace. While residual pathways and stochastic depth appear implicated in block recurrence, isolating causal mechanisms will require controlled training-dynamics at scale; and although two tied blocks recover most of DINOv2, a small residual gap remains that may call for improved recurrent distillation or additional time-varying components. Overall, our work highlights a recurrence induced simplicity bias, suggesting current models admit a recurrent version, implicating a potential simpler analysis. Taken together, this recurrence-induced simplicity bias and its interpretability potential point toward a broader principle: in deep learning, *recurrence finds a way*.

## 6 ACKNOWLEDGMENTS

The authors would like to thank Chris Hamblin, Louis Béthune and John Hammond for many fruitful discussions. This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University. MJ is supported by the Kempner Institute Graduate Research Fellowship. TAK, TF, and RH are supported by the Kempner Institute Research Fellowship. DB is supported by the Kempner Institute for the Study of Natural and Artificial Intelligence at Harvard University.

## REFERENCES

- Julius Adebayo, Michael Muelly, Ilaria Liccardi, and Been Kim. Debugging tests for model explanations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Lucian Alecu, Hugues Bonnin, Thomas Fel, Laurent Gardes, Sébastien Gerchinovitz, Ludovic Ponsolle, Franck Mamalet, Éric Jenn, Vincent Mussot, Cyril Cappi, et al. Can we reconcile safety objectives with machine learning performances? *ERTS*, 2022.
- Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. *ArXiv e-print*, 2021.
- Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *Public Library of Science (PloS One)*, 2015.
- Federico Baldassarre, Marc Szafraniec, Basile Terver, Vasil Khalidov, Francisco Massa, Yann LeCun, Patrick Labatut, Maximilian Seitzer, and Piotr Bojanowski. Back to the features: Dino as a foundation for video world models. *ArXiv e-print*, 2025.
- Randall Balestriero et al. A spline theory of deep learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety—a review. *ArXiv e-print*, 2024.
- Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021.
- Alexandre Boulch. Sharesnet: reducing residual network parameter number by sharing weights. *ArXiv e-print*, 2017.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023.
- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 2017.
- Diogo V. Carvalho, Eduardo M. Pereira, and Jaime S. Cardoso. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 2019.
- Gregory J Chaitin. Algorithmic information theory. *IBM journal of research and development*, 1977.
- Mathieu Chalvidal, Matthew Ricci, Rufin VanRullen, and Thomas Serre. Go with the flow: Adaptive control for neural odes. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Christopher Chiu, Maximilian Heil, Teresa Kim, and Anthony Miyaguchi. Fine-grained classification for poisonous fungi identification with transfer learning. *ArXiv e-print*, 2024.
- Radoslaw M Cichy and Daniel Kaiser. Deep neural networks as scientific models. *Trends in cognitive sciences*, 2019.
- Taco Cohen and Max Welling. Learning the irreducible representations of commutative lie groups, 2014.
- Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *ArXiv e-print*, 2023.
- Sayantan Dasgupta and Trevor Cohn. Improving language model distillation through hidden state matching. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.
- Soham De, Samuel L Smith, Anushan Fernando, Aleksandar Botev, George Cristian-Muraru, Albert Gu, Ruba Haroun, Leonard Berrada, Yutian Chen, Srivatsan Srinivasan, et al. Griffin: Mixing gated linear recurrences with local attention for efficient language models. *ArXiv e-print*, 2024.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *ArXiv e-print*, 2018.
- Kamaludin Dingle, Guillermo Valle Pérez, and Ard A Louis. Generic predictions of output probability based on complexities of inputs and outputs. *Scientific reports*, 2020.
- Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *ArXiv e-print*, 2017.
- 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. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.
- Thomas Fel. *Sparks of explainability: recent advancements in explaining large vision models*. PhD thesis, Université de Toulouse, 2024.
- Thomas Fel, Victor Boutin, Mazda Moayeri, Remi Cadene, Louis Bethune, Mathieu Chalvidal, and Thomas Serre. A holistic approach to unifying automatic concept extraction and concept importance estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2023.
- Thomas Fel, Ekdeep Singh Lubana, Jacob S Prince, Matthew Kowal, Victor Boutin, Isabel Papadimitriou, Binxu Wang, Martin Wattenberg, Demba Ba, and Talia Konkle. Archetypal sae: Adaptive and stable dictionary learning for concept extraction in large vision models. *Proceedings of the International Conference on Machine Learning (ICML)*, 2025.
- Jesseba Fernando and Grigori Guitchounts. Transformer dynamics: A neuroscientific approach to interpretability of large language models. *ArXiv e-print*, 2025.
- Ruth C. Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach, 2025.
- Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. The emergence of clusters in self-attention dynamics. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

- Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- Leilani H. Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. Explaining explanations: An overview of interpretability of machine learning. *Proceedings of the IEEE International Conference on data science and advanced analytics (DSAA)*, 2018.
- Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 1995.
- Micah Goldblum, Marc Finzi, Keefer Rowan, and Andrew Gordon Wilson. The no free lunch theorem, kolmogorov complexity, and the role of inductive biases in machine learning. *ArXiv e-print*, 2023.
- Liv Gorton. The missing curve detectors of inceptionv1: Applying sparse autoencoders to inceptionv1 early vision. *ArXiv e-print*, 2024.
- Klaus Greff, Rupesh K Srivastava, and Jürgen Schmidhuber. Highway and residual networks learn unrolled iterative estimation. *ArXiv e-print*, 2016.
- Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 2017.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *ArXiv e-print*, 2021.
- Anna Hedström, Leander Weber, Dilyara Bareeva, Franz Motzkus, Wojciech Samek, Sebastian Lapuschkin, and Marina M-C Höhne. Quantus: an explainable ai toolkit for responsible evaluation of neural network explanations. *The Journal of Machine Learning Research (JMLR)*, 2022.
- Nhat M. Hoang, Do Xuan Long, Cong-Duy Nguyen, Min-Yen Kan, and Luu Anh Tuan. A comparative analysis of contextual representation flow in state-space and transformer architectures. *ArXiv e-print*, 2025.
- Ann Huang, Mitchell Ostrow, Satpreet H Singh, Leo Kozachkov, Ila Fiete, and Kanaka Rajan. Inputdsa: Demixing then comparing recurrent and externally driven dynamics. *ArXiv e-print*, 2025.
- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Weinberger. Deep networks with stochastic depth, 2016.
- Minyoung Huh, Hossein Mobahi, Richard Zhang, Brian Cheung, Pulkit Agrawal, and Phillip Isola. The low-rank simplicity bias in deep networks. *The Journal of Transactions on Machine Learning Research (TMLR)*, 2023.
- Alexia Jolicoeur-Martineau. Less is more: Recursive reasoning with tiny networks. *ArXiv e-print*, 2025.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models, 2020.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *ArXiv e-print*, 2023.
- Andrei N Kolmogorov. Three approaches to the quantitative definition of information'. *Problems of information transmission*, 1965.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited, 2019.

- Matthew Kowal, Richard P Wildes, and Konstantinos G Derpanis. Visual concept connectome (vcc): Open world concept discovery and their interlayer connections in deep models. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- Leonid Anatolevich Levin. Universal sequential search problems. *Problemy peredachi informatsii*, 1973.
- Ming Li, Paul Vitányi, et al. *An introduction to Kolmogorov complexity and its applications*. Springer, 2008.
- Qianli Liao and Tomaso Poggio. Bridging the gaps between residual learning, recurrent neural networks and visual cortex, 2020.
- Drew Linsley, Alekh Karkada Ashok, Lakshmi Narasimhan Govindarajan, Rex Liu, and Thomas Serre. Stable and expressive recurrent vision models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*. Springer, 2024.
- Max Losch, Mario Fritz, and Bernt Schiele. Semantic bottlenecks: Quantifying and improving inspectability of deep representations. *International Journal of Computer Vision*, 2021.
- Yiping Lu, Zhuohan Li, Di He, Zhiqing Sun, Bin Dong, Tao Qin, Liwei Wang, and Tie-Yan Liu. Understanding and improving transformer from a multi-particle dynamic system point of view. *ArXiv e-print*, 2019.
- Yi Ma, Doris Tsao, and Heung-Yeung Shum. On the principles of parsimony and self-consistency for the emergence of intelligence. *Frontiers of Information Technology & Electronic Engineering*, 2022.
- Chris Mingard, Henry Rees, Guillermo Valle-Pérez, and Ard A Louis. Deep neural networks have an inbuilt occam’s razor. *Nature Communications*, 2025.
- Guido Montúfar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- Sabine Muzellec, Leo Andeol, Thomas Fel, Rufin VanRullen, and Thomas Serre. Gradient strikes back: How filtering out high frequencies improves explanations. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.
- Thao Nguyen, Maithra Raghu, and Simon Kornblith. On the origins of the block structure phenomenon in neural network representations. *arXiv preprint arXiv:2202.07184*, 2022.
- Paul Novello, Thomas Fel, and David Vigouroux. Making sense of dependence: Efficient black-box explanations using dependence measure. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Chris Olah, Nick Cammarata, Chelsea Voss, Ludwig Schubert, and Gabriel Goh. Naturally occurring equivariance in neural networks. *Distill*, 2020.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *ArXiv e-print*, 2023.
- Mitchell Ostrow, Adam Eisen, Leo Kozachkov, and Ila Fiete. Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. *Proceedings of the International Conference on Machine Learning (ICML)*, 2013.

- William Peebles and Saining Xie. Scalable diffusion models with transformers. *ArXiv e-print*, 2023.
- Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *Proceedings of the British Machine Vision Conference (BMVC)*, 2018.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *ArXiv e-print*, 2021.
- Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in neural information processing systems*, 34:12116–12128, 2021.
- Michael Sander, Pierre Ablin, and Gabriel Peyré. Do residual neural networks discretize neural ordinary differential equations? *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv e-print*, 2019.
- Peter J Schmid. Dynamic mode decomposition and its variants. *Annual Review of Fluid Mechanics*, 2022.
- Jürgen Schmidhuber. Discovering neural nets with low kolmogorov complexity and high generalization capability. *Neural Networks*, 1997.
- Jürgen Schmidhuber. The speed prior: a new simplicity measure yielding near-optimal computable predictions. *International conference on computational learning theory*, 2002.
- Avi Schwarzschild, Arjun Gupta, Amin Ghiasi, Micah Goldblum, and Tom Goldstein. The uncanny similarity of recurrence and depth. *arXiv preprint arXiv:2102.11011*, 2021.
- Thiago Serra, Christian Tjandraatmadja, and Srikumar Ramalingam. Bounding and counting linear regions of deep neural networks. *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Peter Shaw, James Cohan, Jacob Eisenstein, and Kristina Toutanova. Bridging kolmogorov complexity and deep learning: Asymptotically optimal description length objectives for transformers. *ArXiv e-print*, 2025.
- Sam Shleifer and Alexander M Rush. Pre-trained summarization distillation. *ArXiv e-print*, 2020.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Shawn Tan, Yikang Shen, Zhenfang Chen, Aaron Courville, and Chuang Gan. Sparse universal transformer. *ArXiv e-print*, 2023.
- Matus Telgarsky. Representation benefits of deep feedforward networks. *ArXiv e-print*, 2015.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024.
- Vincent Tjeng and Russ Tedrake. Verifying neural networks with mixed integer programming. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

- Guillermo Valle-Perez, Chico Q Camargo, and Ard A Louis. Deep learning generalizes because the parameter-function map is biased towards simple functions. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- Andreas Veit, Michael J Wilber, and Serge Belongie. Residual networks behave like ensembles of relatively shallow networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- Siddarth Venkatraman, Vineet Jain, Sarthak Mittal, Vedant Shah, Johan Obando-Ceron, Yoshua Bengio, Brian R Bartoldson, Bhavya Kailkhura, Guillaume Lajoie, Glen Berseth, et al. Recursive self-aggregation unlocks deep thinking in large language models. *ArXiv e-print*, 2025.
- Martina G Vilas, Timothy Schaumlöffel, and Gemma Roig. Analyzing vision transformers for image classification in class embedding space. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Yue Wang and Sai Ho Chung. Artificial intelligence in safety-critical systems: a systematic review. *Industrial Management & Data Systems*, 2022.
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1989.
- Andrew Gordon Wilson. Deep learning is not so mysterious or different. *ArXiv e-print*, 2025.
- Zuxuan Wu, Tushar Nagarajan, Abhishek Kumar, Steven Rennie, Larry S. Davis, Kristen Grauman, and Rogerio Feris. Blockdrop: Dynamic inference paths in residual networks, 2019.
- Qing Yun. Vision transformers (vits) for feature extraction and classification of ai-generated visual designs. *IEEE Access*, 2025.
- Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2014.

## A TRAINING BLOCK RECURRENT FOUNDATION MODELS

All Raptor variants ( $k = 2, 3, 4$ ) are trained on top of DINOv2 (ViT-B with registers) (Darcet et al., 2023) and use the same transformer architecture. The model utilizes a feature dimension of 784, an MLP ratio of 4, and 12 multi-head attention heads. The depth-scale MLP in each block generates three separate scaling vectors; we provide the full formulation in Appendix A.3. For each Raptor variant in Table 2 and Figure 5, we train three models, each with a different seed. We measure the performance of each model and report aggregate statistics. Below, we provide the hyperparameters and training settings used to train Raptor, describing the layer divisions for different values of  $k$  followed by the training procedure.

For each choice of  $k$ , the encoder layers are divided into blocks as shown in Table 3. We select these partitions using the max cut algorithm applied to the DINOv2 ViT-B activation layer-layer cosine similarity matrix using 10000 samples from the ImageNet-1k training set.

$k$	Block	Input $\rightarrow$ Predicted Layers
2	Block 1	0 $\rightarrow$ 1–7
	Block 2	7 $\rightarrow$ 8–12
3	Block 1	0 $\rightarrow$ 1–7
	Block 2	7 $\rightarrow$ 8–10
	Block 3	10 $\rightarrow$ 11–12
4	Block 1	0 $\rightarrow$ 1–4
	Block 2	4 $\rightarrow$ 5–7
	Block 3	7 $\rightarrow$ 8–10
	Block 4	10 $\rightarrow$ 11–12

Table 3: Layer splits used for training with different values of  $k$ .

In the first stage, each block is trained independently on the subset of layers it is responsible for predicting. For example, when  $k = 3$ , Block 1 is trained to predict Layers 1–7. We train on the ImageNet-1k train split for 20 epochs with a batch size of 64 using the AdamW optimizer with a weight decay of 0.0001. The learning rate follows a linear warmup for 10,000 steps to  $1 \times 10^{-4}$ , followed by a cosine decay to  $1 \times 10^{-6}$ . The teacher forcing loss weight  $\lambda$  is annealed from 0.5 to 0 over the first 5 epochs. For the third block, we use specific token loss weights of  $\lambda_{\text{cls}} = 0.34$ ,  $\lambda_{\text{reg}} = 0.33$ , and  $\lambda_{\text{patch}} = 0.33$ .

In the second stage, after independent training, all blocks are connected to autoregressively predict Layers 1–12 end-to-end. Each block still predicts its designated segment, but the entire model now backpropagates through the full sequence. We maintain the same dataset, epoch count, batch size, weight decay, optimizer, and learning rate schedule as the first stage. However, the token loss weights are adjusted to  $\lambda_{\text{cls}} = 0.45$ ,  $\lambda_{\text{reg}} = 0.1$ , and  $\lambda_{\text{patch}} = 0.45$ .

The number of recurrent timesteps each block executes corresponds directly to its layer coverage; for instance, in the  $k=3$  setup, Block 1 runs for 7 timesteps, Block 2 for 3 timesteps, and Block 3 for 2 timesteps.

For a given batch of ground truth DINOv2 activations  $\mathbf{X} \in \mathbb{R}^{N \times T \times D}$  and Raptor predictions  $\hat{\mathbf{X}} \in \mathbb{R}^{N \times T \times D}$ , where  $N$  denotes the batch size,  $T$  the number of tokens, and  $D$  the feature dimension.

We decompose the activation tensor along the token dimension into three distinct components corresponding to the token types: the class token  $\mathbf{X}_{\text{cls}} \in \mathbb{R}^{N \times 1 \times D}$ , the register tokens  $\mathbf{X}_{\text{reg}} \in \mathbb{R}^{N \times N_{\text{reg}} \times D}$ , and the spatial patch tokens  $\mathbf{X}_{\text{patch}} \in \mathbb{R}^{N \times N_{\text{patch}} \times D}$ , such that  $\mathbf{X} = \text{concat}(\mathbf{X}_{\text{cls}}, \mathbf{X}_{\text{reg}}, \mathbf{X}_{\text{patch}})$ .

We define the total reconstruction loss function  $\mathcal{L}(\mathbf{X}, \hat{\mathbf{X}})$  as a weighted sum of the errors for each component:

$$\begin{aligned} \mathcal{L}(\mathbf{X}, \hat{\mathbf{X}}) = & \lambda_{\text{cls}} \|\hat{\mathbf{X}}_{\text{cls}} - \mathbf{X}_{\text{cls}}\|_F^2 + \lambda_{\text{reg}} \|\hat{\mathbf{X}}_{\text{reg}} - \mathbf{X}_{\text{reg}}\|_F^2 \\ & + \lambda_{\text{patch}} \|\hat{\mathbf{X}}_{\text{patch}} - \mathbf{X}_{\text{patch}}\|_F^2, \end{aligned} \quad (3)$$

### A.1 PHASE DISCOVERY VIA A CONTIGUOUS MAX-CUT ON THE LAYER-LAYER SIMILARITY

**Problem setup.** Let  $S \in \mathbb{R}^{L \times L}$  be the (symmetrized) layer-layer similarity matrix, where  $S_{ij}$  measures the similarity between layers  $i$  and  $j$  (for example, cosine similarity). We seek a partition of depth into  $k$  contiguous segments or “phases”.  $\Pi = \{[b_1, e_1], \dots, [b_k, e_k]\}$  with  $1 = b_1 \leq e_1 < b_2 \leq e_2 < \dots < b_k \leq e_k = L$  and  $e_t + 1 = b_{t+1}$ , that maximizes within-block similarity (equivalently, minimizes cross-block cut).

**Objectives.** For a segment  $[i, j]$  of length  $n = j - i + 1$ , define:

$$\text{sum}(i, j) = \sum_{p=i}^j \sum_{q=i}^j S_{pq}, \quad \text{offdiag}(i, j) = \text{sum}(i, j) - \sum_{p=i}^j S_{pp}.$$

We consider additive segment scores  $g(i, j)$  by computing the final weighted mean as:

$$\frac{\text{sum}(i, j)}{n^2};$$

Maximizing  $\sum_{t=1}^k g(b_t, e_t)$  prefers blocks that are internally similar and, by contiguity, implies small cross-block interfaces (a contiguous max-cut on the line).

**Fast block queries via 2-D prefix sums.** Precompute a 2-D prefix (summed-area) table  $P \in \mathbb{R}^{(L+1) \times (L+1)}$  with  $P_{rc} = \sum_{u < r} \sum_{v < c} S_{uv}$ . Then any submatrix sum obeys

$$\text{sum}(i, j) = P_{j+1, j+1} - P_{i, j+1} - P_{j+1, i} + P_{i, i},$$

in  $O(1)$  time; diagonal sums use a 1-D prefix over  $\text{diag}(S)$ . This is sometimes referred to as the “integral image” trick.

**Contiguous DP solver ( $O(kL^2)$ ).** Let  $\text{dp}[t, j]$  be the best score for partitioning layers  $1..j$  into  $t$  blocks. With minimum block length  $m$ ,

$$\text{dp}[1, j] = g(1, j) \quad (j \geq m), \quad \text{dp}[t, j] = \max_{i \in \{t m - 1, \dots, j - m\}} \text{dp}[t - 1, i] + g(i + 1, j),$$

for  $t = 2, \dots, k$  and  $j \geq t m$ . We keep backpointers  $\text{prev}[t, j]$  to recover boundaries by backtracking from  $(t=k, j=L)$ . With  $g(\cdot)$  evaluated in  $O(1)$  by prefix sums, the overall complexity is  $O(kL^2)$  time and  $O(kL)$  memory. This DP structure mirrors classical optimal 1-D segmentation/partitioning solvers.

### A.2 TEACHER-STUDENT RECONSTRUCTION $R^2$

To stabilize measures of pairwise vector similarity over large hyperparameter sweeps when fits may be poor, we use an alternative calculation for  $R^2$  for small-scale models (Figure 4). Here, we first regress the student’s c1s or patch tokens  $\in \mathbb{R}^{N \times D}$  to the corresponding teacher tokens  $\in \mathbb{R}^{N \times D}$  using ordinary least squares with a bias term.  $N$  is the number of tokens and  $D$  is the dimensionality of the token. We then calculate the average of the  $R^2$  values between the true teacher token vectors and the student’s reconstruction of those vectors. Note that this regression is purely a 1-dimensional rescaling and shifting for each student token vector. This results in an  $R^2$  value that is bounded between 0 and 1, and can be understood to present the ‘explainable variance’ between the student and teacher representations.

### A.3 DEPTH-SCALING MECHANISM

We employ a lightweight conditioning MLP in each block that maps a scalar depth coordinate  $z \in \mathbb{R}$  to layer-specific scaling vectors. For a layer with feature dimension  $D = 784$ , the depth-scale MLP generates a modulation vector  $\mathbf{v} \in \mathbb{R}^{3D}$ .

The MLP projects the scalar input to a hidden dimension  $h = 16$ , applies a SiLU activation, and projects to the output dimension. We add a unity offset to the MLP output to ensure the scaling factors center around 1 (identity behavior) at initialization. The compound scaling vector  $\mathbf{S}$  is computed as:

$$\mathbf{S} = (\mathbf{W}_2 \cdot \text{SiLU}(\mathbf{W}_1 z + \mathbf{b}_1) + \mathbf{b}_2) + \mathbf{1} \quad (4)$$

where  $\mathbf{1}$  denotes a vector of ones. We initialize the weights and biases of the final projection ( $\mathbf{W}_2, \mathbf{b}_2$ ) with a normal distribution ( $\mu = 0, \sigma = 10^{-4}$ ), ensuring that the block initially acts close to standard behavior.

The vector  $\mathbf{S}$  is split along the feature dimension into three component vectors:  $\mathbf{s}_{\text{attn}}, \mathbf{s}_{\text{mlp}}, \mathbf{s}_{\text{out}} \in \mathbb{R}^D$ . These vectors modulate the signal via element-wise multiplication ( $\odot$ ) at three points in the transformer block: the attention residual, the MLP residual, and the final block output. Given an input sequence  $\mathbf{X}$ , the forward pass is defined as:

$$\mathbf{X}' = \mathbf{X} + \mathbf{s}_{\text{attn}} \odot \text{LS}_1(\text{Attn}(\text{LN}_1(\mathbf{X}))) \quad (5)$$

$$\mathbf{X}'' = \mathbf{X}' + \mathbf{s}_{\text{mlp}} \odot \text{LS}_2(\text{MLP}(\text{LN}_2(\mathbf{X}')))) \quad (6)$$

$$\mathbf{Y} = \mathbf{s}_{\text{out}} \odot \mathbf{X}'' \quad (7)$$

where LS refers to LayerScale, LN is Layer Normalization, and broadcasting is applied over the sequence length dimension.

#### A.4 LINEAR PROBE FINE-TUNING

We fine-tune linear probes on three downstream datasets: ImageNet-1k (classification), ADE20k (semantic segmentation), and NYUv2 (monocular depth estimation).

For ImageNet-1k and ADE20k, we use the AdamW optimizer with linear warmup followed by cosine learning rate decay. For NYUv2, we use AdamW with GradScaler and mixed precision training. All probes operate on the final block’s prediction of Layer 12, using either the c1s token, patch tokens, or both, depending on the task. The detailed hyperparameters are shown in Table 4.

For NYUv2, we adopt an approach similar to [Oquab et al. \(2023\)](#). Specifically, we use images at a  $480 \times 640$  resolution and center pad them so that the dimensions are multiples of 14. We feed the images through the model and extract the predictions from the final layer. The c1s token is concatenated with the patch tokens, and the spatial resolution is upsampled by a factor of 4. Both the c1s and patch tokens are upsampled, after which the c1s token is concatenated to each patch token. We treat this representation as the “logits.” To obtain depth, we normalize the logits with a softmax and compute the weighted average of the centers of 256 evenly spaced bins. Then, we upsample this representation to  $480 \times 640$  and consider the result our depth. For training, we use the loss function introduced by [Bhat et al. \(2021\)](#).

Hyperparameter	ImageNet-1k	ADE20k	NYUv2
Epochs	15	10	25
Batch Size	512	32	128
Base LR	$1 \times 10^{-3}$	$1 \times 10^{-3}$	$1 \times 10^{-4}$
Weight Decay	$1 \times 10^{-2}$	$1 \times 10^{-2}$	$1 \times 10^{-2}$
Grad. Clip Norm	1.0	1.0	1.0
Warmup Iters	100	100	100
Optimizer	AdamW	AdamW	AdamW + GradScaler
Head Init.	DINOv2 classification probe	Random segmentation head	Random depth head
Input Tokens	concat(c1s, mean patch)	Patch	concat(c1s, patch)

Table 4: Linear probe fine-tuning hyperparameters across datasets. Base LR denotes the peak learning rate before cosine decay.

## B EVALUATING MAX-CUT ALGORITHM

To evaluate the efficacy of the max cut algorithm (Figure 3), we initialize a ViT-B with ImageNet-21k and ImageNet-1k weights and replace the classifier with a randomized linear probe applied

to the CLS token. We fine-tune the full model on CIFAR-100 for 10 epochs, achieving 90.7% validation accuracy. Using activations extracted from the CIFAR-100 training set, we compute a layer-wise cosine similarity matrix and apply our max cut algorithm to identify Raptor partitions for  $k \in \{2, 3, 4\}$ . We compare these against two baselines: 10 random contiguous partitions and 10 random non-contiguous partitions (labeled ‘‘Random Shuffle’’), ensuring the max-cut solution is excluded from the samples. Finally, we train Raptor models on each partition configuration for 100 epochs and evaluate performance using the linear probe from the initial fine-tuning stage.

## C DYNAMICS PROTOCOLS AND METRICS

This appendix consolidates definitions and experimental procedures used in Sec. 4. All measurements are performed on ImageNet validation data. For aggregate statistics, we use 10k randomly sampled validation images. For trajectory visualizations (e.g., Fig. 6), we select five ImageNet classes with 1,000 images each. Inputs are resized to 256 pixels on the shorter side and center-cropped to  $224 \times 224$ . Unless otherwise noted, we use DINOv2-Giant with four register tokens from the official implementation.

**Normalization.** Token states  $\mathbf{x}_\ell \in \mathbb{R}^d$  at depth  $\ell$  are decomposed into norm and direction. We study normalized states

$$\hat{\mathbf{x}}_\ell = \frac{\mathbf{x}_\ell}{\|\mathbf{x}_\ell\|} \in \mathbb{S}^{d-1},$$

so that dynamics are restricted to the unit sphere.

**Directional convergence.** Directional similarity to the terminal representation is measured by

$$\gamma_\ell = \langle \hat{\mathbf{x}}_\ell, \hat{\mathbf{x}}_L \rangle,$$

which traces the angular alignment of layer  $\ell$  to the final state.

**Angular speed.** Per-layer angular update magnitude is defined as

$$s_\ell = \arccos \langle \hat{\mathbf{x}}_{\ell+1}, \hat{\mathbf{x}}_\ell \rangle.$$

Statistics of  $s_\ell$  are stratified by token type.

**Phase overlays.** Phase boundaries are obtained from the max-cut segmentation of representational similarity matrices (Sec. 2) and used as vertical markers in angular speed and sensitivity plots.

**Perturbation protocol.** To probe stability, we add a perturbation  $\varepsilon \mathbf{u}$  at layer  $\ell$ ,

$$\tilde{\mathbf{x}}_\ell = \mathbf{x}_\ell + \varepsilon \mathbf{u}, \quad \mathbf{u} \sim \mathcal{N}(0, I_d),$$

and follow the normalized trajectory thereafter. Sensitivity is quantified by the terminal cosine deviation

$$d_{\cos}(\hat{\mathbf{x}}_L^{(\varepsilon, \ell)}, \hat{\mathbf{x}}_L) = 1 - \langle \hat{\mathbf{x}}_L^{(\varepsilon, \ell)}, \hat{\mathbf{x}}_L \rangle.$$

**Low-rank and coherence metrics.** For angular updates  $\Delta_\ell^{(i)} = \hat{\mathbf{x}}_{\ell+1}^{(i)} - \hat{\mathbf{x}}_\ell^{(i)}$ , we form the update matrix  $\mathbf{U}_\ell$ . Stable rank is given by

$$r_s(\mathbf{U}_\ell) = \frac{\|\mathbf{U}_\ell\|_F^2}{\|\mathbf{U}_\ell\|_2^2},$$

and coherence by

$$\kappa_\ell = \frac{1}{N} \sum_i \frac{\langle \Delta_\ell^{(i)}, \bar{\Delta}_\ell \rangle}{\|\Delta_\ell^{(i)}\| \|\bar{\Delta}_\ell\|}, \quad \bar{\Delta}_\ell = \frac{1}{N} \sum_i \Delta_\ell^{(i)}.$$

## D DYNAMIC MODE DECOMPOSITION

Let  $\mathbf{f}$  be a trained ViT with transformer layers  $\{\mathbf{f}_\ell\}_{\ell=1}^L$ . For  $\mathbf{x} \in \mathcal{X}$ , denote by  $\mathbf{A}_\ell(\mathbf{x}) \in \mathbb{R}^{T \times d}$  the token matrix at depth  $\ell$  with  $T = 1 + R + P$  (cls,  $R$  registers,  $P$  patch). Form group states by

within-layer averaging

$$\mathbf{z}_\ell^{(\text{cls})}(x) = \mathbf{A}_\ell(x)_{\text{cls}} \quad \mathbf{z}_\ell^{(\text{reg})}(x) = \frac{1}{R} \sum_{t \in \mathcal{T}_{\text{reg}}} \mathbf{A}_\ell(x)_t \quad \mathbf{z}_\ell^{(\text{patch})}(x) = \frac{1}{P} \sum_{t \in \mathcal{T}_{\text{patch}}} \mathbf{A}_\ell(x)_t$$

and enforce per-layer  $\ell_2$  normalization on the group averages

$$\mathbf{x}_\ell^{(g)}(x) = \frac{\mathbf{z}_\ell^{(g)}(x)}{\|\mathbf{z}_\ell^{(g)}(x)\|_2} \in \mathbb{S}^{d-1} \subset \mathbb{R}^d.$$

All DMD fits below are performed independently for each  $g \in \{\text{cls}, \text{reg}, \text{patch}\}$  on the depth trajectory  $\mathbf{x}_{0:L}^{(g)}(x)$ . We start by stacking states along depth to form

$$\mathbf{Y}^{(g)} = \begin{bmatrix} (\mathbf{x}_0^{(g)})^\top \\ \vdots \\ (\mathbf{x}_L^{(g)})^\top \end{bmatrix} \in \mathbb{R}^{(L+1) \times d} \quad \mathbf{X}_1 = (\mathbf{Y}_{0:L}^{(g)})^\top \in \mathbb{R}^{d \times L} \quad \mathbf{X}_2 = (\mathbf{Y}_{1:L+1}^{(g)})^\top \in \mathbb{R}^{d \times L}.$$

DMD fits a single linear depth-step map  $\mathbf{A}$  with  $\mathbf{X}_2 \approx \mathbf{A}\mathbf{X}_1$ . For the exact DMD at rank  $r$ , we compute the SVD  $\mathbf{X}_1 = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  and select  $r \leq \text{rank}(\mathbf{X}_1)$ . Let  $\mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r$  be the leading blocks and define the reduced operator

$$\tilde{\mathbf{A}} = \mathbf{U}_r^\top \mathbf{X}_2 \mathbf{V}_r \mathbf{\Sigma}_r^{-1} \in \mathbb{R}^{r \times r}.$$

Diagonalize  $\tilde{\mathbf{A}}\mathbf{W} = \mathbf{W}\mathbf{\Lambda}$  with  $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_r) \in \mathbb{C}^{r \times r}$  and  $\mathbf{W} \in \mathbb{C}^{r \times r}$ . The exact DMD modes in ambient space are

$$\mathbf{\Phi} = \mathbf{X}_2 \mathbf{V}_r \mathbf{\Sigma}_r^{-1} \mathbf{W} \in \mathbb{C}^{d \times r}.$$

Modal amplitudes for the initial state are  $\mathbf{b} = \mathbf{\Phi}^\dagger \mathbf{x}_0^{(g)} \in \mathbb{C}^r$ . One-step and  $t$ -step reconstructions are

$$\hat{\mathbf{x}}_1^{(g)} \approx \mathbf{\Phi}\mathbf{\Lambda}\mathbf{b} \quad \hat{\mathbf{x}}_t^{(g)} \approx \mathbf{\Phi}\mathbf{\Lambda}^t\mathbf{b} \text{ for } t \geq 0$$

with the induced linear predictor  $\mathbf{A} = \mathbf{X}_2 \mathbf{V}_r \mathbf{\Sigma}_r^{-1} \mathbf{U}_r^\top$ . No affine offset is fitted.

**On-sphere interpretation.** Because each  $\mathbf{x}_\ell^{(g)}$  lies on  $\mathbb{S}^{d-1}$  the map  $\mathbf{A}$  is a best linear approximation of the depth flow restricted to the unit sphere. For  $\lambda_i = |\lambda_i|e^{i\theta_i}$  the modulus  $|\lambda_i|$  measures contraction of directions within the span of modes on the sphere and the angle  $\theta_i$  captures per-layer rotational change. The spectral radius  $\rho(\mathbf{A}) = \max_i |\lambda_i|$  and the median of  $|\lambda_i|$  summarize contraction strength. In particular this explain why all the eigenvalue in [Figure 9](#) are contain in  $\mathbb{S}^2$ , see the report of the eigenvalue cloud  $\{\lambda_i\}_{i=1}^r$  in the complex plane and the singular spectrum  $\{\sigma_i\}_{i=1}^r$  of  $\mathbf{X}_1$ .

## E LEVIN COMPLEXITY UNDER 0-BRH

**Background & Positioning.** There is a long line of work linking plain Kolmogorov complexity [Kolmogorov \(1965\)](#); [Chaitin \(1977\)](#) to its resource-bounded variants, notably Levin’s time-bounded measure, which penalizes programs by both description length and runtime (the  $|p| + \log T$  form) ([Levin, 1973](#); [Li et al., 2008](#)). Deep learning research has not remained isolated from these ideas: connections have been explored both for discovering simpler neural networks [Schmidhuber \(1997\)](#) and for contextualizing runtime priors [Schmidhuber \(2002\)](#) within neural architectures.

In our case, the proof is a direct application of this standard framework as formalized in [Li et al. \(2008\)](#): once we exhibit a prefix-free program  $p$  that computes  $\mathbf{f}_\ell$  and bound its running time  $T(p)$ , Levin complexity immediately yields  $K_{\text{Levin}}(\mathbf{f}_\ell) \leq |p| + \log T(p) + O(1)$ . The only domain-specific ingredients are (i) a BRH-based encoding that prices the tied blocks and schedule at  $\sum_{j=1}^k DL_U(\boldsymbol{\theta}(\mathbf{B}_j)) + O(k \log \ell)$  using standard self-delimiting integer codes, and (ii) a runtime-parity assumption linking  $T(p)$  to the model’s untied runtime  $R(\mathbf{f}_\ell)$  up to machine and invariance constants. Formally:

**Claim 2** (BRH guarantees low Levin complexity). *Let  $\mathbf{f}_\ell$  satisfy 0-BRH with  $k \ll \ell$  tied blocks  $\{\mathbf{B}_j\}_{j=1}^k$  and schedule  $(n_j)_{j=1}^k$  with  $\sum_j n_j = \ell$ . Let  $R(\cdot)$  denote block runtime and define the untied*

teacher runtime  $R(\mathbf{f}_\ell) := \sum_{j=1}^k n_j R(\mathbf{B}_j)$ ; assume runtime parity  $R(\mathbf{B}_j) \leq (1 + \delta)R(\mathbf{B}_f)$ . Then

$$K_{\text{Levin}}(\mathbf{f}_\ell) \leq \sum_{j=1}^k DL(\boldsymbol{\theta}(\mathbf{B}_j)) + O(k \log \ell) + \log R(\mathbf{f}_\ell) + O(1).$$

*Proof.* We follow the standard time-bounded description framework of Levin (1973), and contextualized it to the 0-BRH setting (perfect reconstruction along the entire depth trajectory), in the spirit of the sketch already outlined in the paper. We fix a universal prefix-free machine  $U$  and adopt the same finite-precision arithmetic used by the target model to define  $DLU(\cdot)$ ; this ensures that encoding parameters with length  $DLU(\boldsymbol{\theta}(\mathbf{B}_j))$  bits suffices to reproduce the exact numerical behavior of the corresponding block under that arithmetic. With this convention, “exactly computes” means bit-for-bit equivalence with the reference implementation of  $\mathbf{f}_\ell$  (under its inference precision).

The proof proceeds in three steps: we first bound the description length, then verify correctness (equality of the program output), and finally bound the runtime before combining these results to derive the final bound.

The first step consists in bounding the description length. We describe a prefix-free program  $p$  that parses (i) the tied block parameters and (ii) a self-delimiting schedule, with total length

$$|p| \leq \sum_{j=1}^k DLU(\boldsymbol{\theta}(\mathbf{B}_j)) + O(k \log \ell) + O(1).$$

For each tied block, encode its parameters using the scheme underlying  $DLU$ , yielding a prefix-free code of length  $DLU(\boldsymbol{\theta}(\mathbf{B}_j))$ . Concatenate the  $k$  codes and prepend an  $O(1)$  header that instructs  $U$  how to parse and reconstruct each  $\mathbf{B}_j$  from its bitstream. Encode the schedule  $(n_1, \dots, n_k)$  and the value of  $k$  as self-delimiting integers (any standard universal code suffices), which costs  $O(k \log \ell)$  bits because  $n_j \leq \ell$  and  $\sum_j n_j = \ell$ . The parser is a constant-size routine bundled with  $p$ , absorbed into  $O(1)$ .

Now we show that this construction is a valid under 0-BRH. Let  $g_\ell := \mathbf{B}_k^{(n_k)} \circ \dots \circ \mathbf{B}_1^{(n_1)}$  denote the composed tied map realized by iterating the reconstructed blocks according to the parsed schedule. By 0-BRH, we have exact functional equality

$$\mathbf{f}_\ell = \mathbf{B}_k^{(n_k)} \circ \dots \circ \mathbf{B}_1^{(n_1)} = g_\ell,$$

i.e., the tied computation reproduces the full internal trajectory (not only the terminal output). Because  $p$  encodes the same block parameters at the precision that defines  $DLU$  and  $U$  evaluates precisely the same composition, we obtain, for all inputs  $x$ ,

$$U(p)(x) = g_\ell(x) = \mathbf{f}_\ell(x).$$

Finally, we get the runtime bound to derive the Levin Complexity. On the fixed machine  $U$ , a single evaluation of block  $\mathbf{B}_j$  incurs time at most  $cR(\mathbf{B}_j)$  for a machine dependent constant  $c > 0$  (standard machine invariance up to constants; cf. the invariance discussion in Kolmogorov/Levin-style arguments (Kolmogorov, 1965; Levin, 1973)). Executing the schedule therefore costs

$$T(p) \leq c \sum_{j=1}^k n_j R(\mathbf{B}_j).$$

Applying the runtime parity assumption  $R(\mathbf{B}_j) \leq (1 + \delta)R(\mathbf{B}_f)$  for all  $j$  and using  $R(\mathbf{f}_\ell) := \sum_{j=1}^k n_j R(\mathbf{B}_j)$  gives

$$T(p) \leq c(1 + \delta) \sum_{j=1}^k n_j R(\mathbf{B}_f) = c(1 + \delta) R(\mathbf{f}_\ell).$$

By the definition of Levin’s time-bounded complexity

$$K_{\text{Levin}}(\mathbf{f}_\ell) \leq |p| + \log T(p) + O(1),$$

and combining with the bounds from the previous steps yields

$$K_{\text{Levin}}(\mathbf{f}_\ell) \leq \sum_{j=1}^k \text{DLU}(\boldsymbol{\theta}(\mathbf{B}_j)) + O(k \log \ell) + \log R(\mathbf{f}_\ell) + \underbrace{\log c + \log(1 + \delta) + O(1)}_{= O(1)}.$$

Absorbing the machine constant and the parity slack into  $O(1)$  gives the stated inequality.  $\square$

## F TRAINING LOSSES AND ALGORITHM

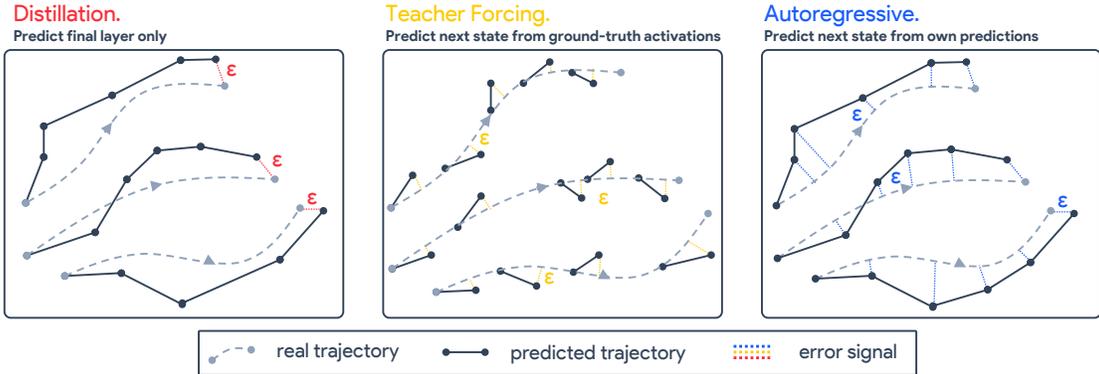


Figure 10: **Three training paradigms for learning recurrent approximations.** Each panel shows three token trajectories through depth. Gray dashed lines with filled circles represent the ground-truth teacher trajectories; black solid lines with filled circles show the student’s predictions; colored dotted lines (with  $\epsilon$  labels) indicate the error signal between predicted and ground-truth states. **Left (Distillation):** The student network directly predicts the final layer from the initial state, with no supervision on intermediate representations. Error is measured only at the terminal state, providing no guidance on the representational trajectory. **Middle (Teacher Forcing):** At each depth step  $\ell$ , the student block predicts  $\hat{x}_{\ell+1} = B(x_\ell)$  using the ground-truth activation  $x_\ell$  from the teacher. Vertical arrows indicate where the student “resets” to ground-truth states. This enables efficient parallel training and prevents error accumulation, but creates a train-test mismatch since the model never learns to handle its own prediction errors. **Right (Autoregressive):** The student autoregressively predicts  $\hat{x}_{\ell+1} = B(\hat{x}_\ell)$  using its own previous predictions, matching inference conditions. Errors can compound across depth (shown by increasing deviation between trajectories), requiring the model to learn self-consistent, closed-loop dynamics. Our two-stage training (Sec. 3) combines both approaches: Stage 1 uses teacher forcing for stable, parallelizable pretraining; Stage 2 switches to autoregressive training to ensure self-consistency at inference.

```

a_gt = teacher(batch) # extract teacher activations
a_ar = a_gt[0] # start at first hidden state
loss = 0

for l in range(num_layers):
    b = get_block(l) # get recurrent block for phase
    out_tf, out_ar = b(a_gt[l]), b(a_ar) # dual forward passes

    # explicit hybrid loss
    l_tf = dist(out_tf, a_gt[l+1])
    l_ar = dist(out_ar, a_gt[l+1])
    loss += lambda * l_tf + (1 - lambda) * l_ar

    a_ar = out_ar # update student state

loss.backward()
optimizer.step()

```

Figure 11: **Hybrid training loop.** The model minimizes a convex combination of Teacher Forcing and Autoregressive errors, balanced by  $\lambda$ .

## G SUPPLEMENTARY RESULTS

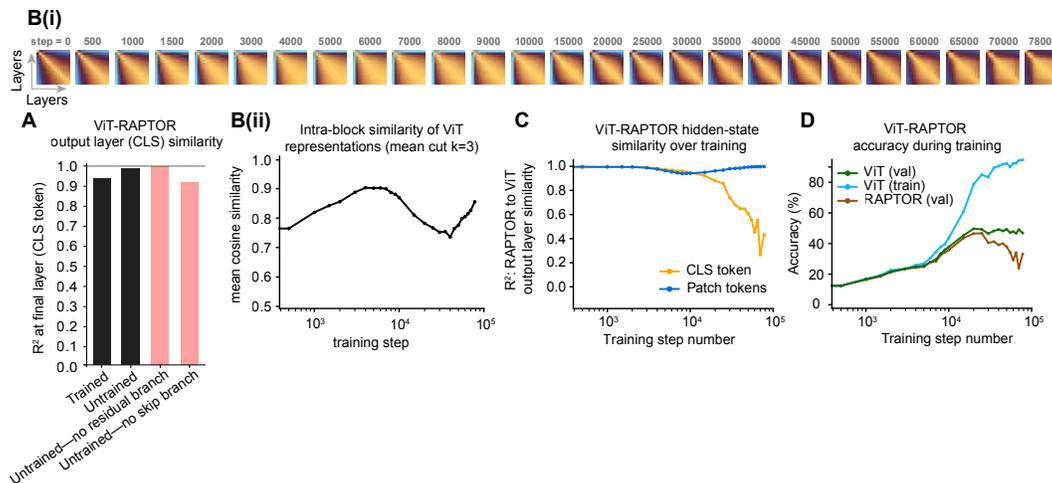


Figure 13: **Similarity and distillation dynamics during training and overtraining.** **A)** ViT-Raptor similarities between four pairs of ViT and Raptor networks. First from the left, a fully trained (but not overtrained) ViT network trained on CIFAR-100 and a Raptor student network trained (with  $k=3$ ). Second, an untrained ViT network, and a Raptor student network. Third, an untrained ViT network with no residual branch in its architecture, so the entire network is just skip connections; and a Raptor student network. Fourth, an untrained ViT network with no skip connections, so the network no longer contains skip connections and is simply a feedforward network; and a Raptor student network. **Bi)** Layer-layer cosine similarities during training in a ViT network that is allowed to be overtrained (overtraining starts at around step 10000, see (D)). **Bii)** Mean intra-block cosine similarity with max-cut  $k=3$  during training. **C)** Teacher-student Raptor reconstruction of ViT during training showing divergence after overtraining starts. **D)** CIFAR-100 accuracy in a ViT that was allowed to overtrain. Divergence between ‘train’ (green) and ‘val’ (blue) lines shows overtraining. X-axis is the same as in panels (B) and (c).

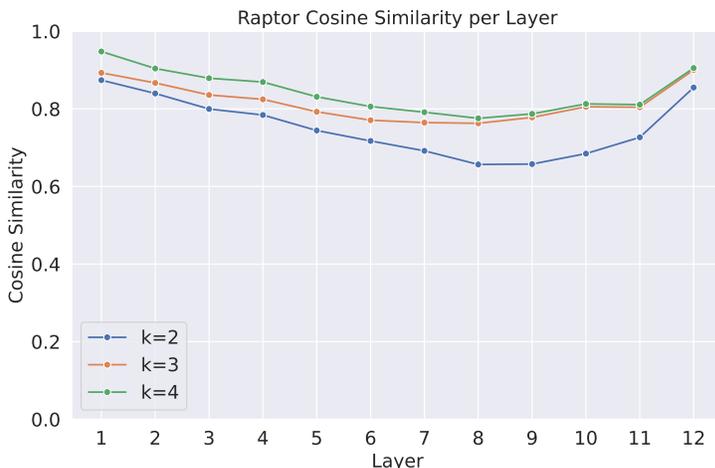


Figure 14: **Cosine similarity remains high through layers.** Cosine similarity between activations from Raptor models ( $k = 2, k = 3, k = 4$ ) and activations from DINOv2-Base on the ImageNet-1k validation set. The consistently high similarity (mostly  $>0.7$ ) indicates that Raptor effectively captures the dynamics of the original model, with the high alignment in the final layer confirming that the block-recurrent mechanism successfully approximates the target output representations.

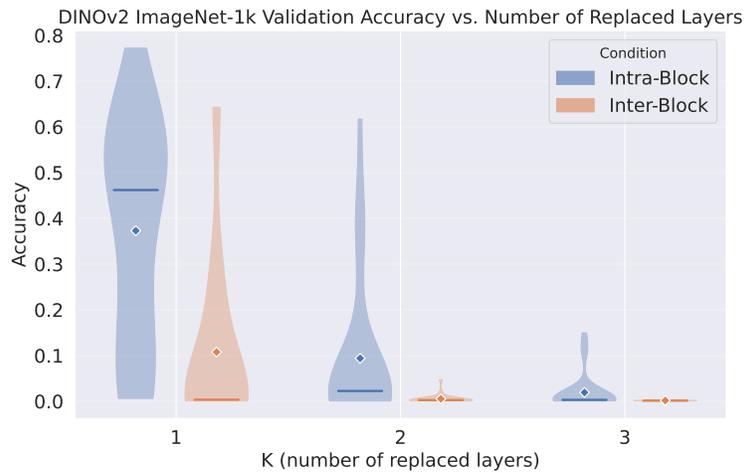


Figure 15: **Causal intervention.** DINOv2-Base accuracy on ImageNet-1k validation set with 1, 2, and 3 layers replaced with another layer ( $k = 1, k = 2, k = 3$ , respectively). Intra-block refers to replacing a layer with another layer from the same block. Inter-block refers to replacing a layer with a layer from a different block. Blocks are determined by the max cut algorithm. The significantly higher accuracy of intra-block replacements (blue) compared to inter-block (orange) confirms that layers within a block are functionally interchangeable in a way that any two arbitrary blocks are not, supporting the block-recurrent hypothesis.

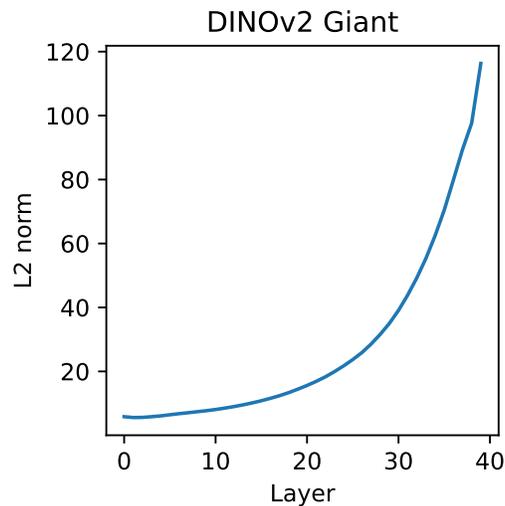


Figure 12: **Depth-wise feature norms.** Magnitudes grow with depth, motivating analysis on directions.