EVIDENCE AGAINST IMPLICITLY RECURRENT COMPUTATIONS IN RESIDUAL NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent work on residual neural networks (ResNets) has suggested that a ResNet's deep feedforward computation may be characterized as implicitly recurrent in that it iteratively refines the same representation like a recurrent network. To test this hypothesis, we manipulate the degree of weight sharing across layers in ResNets using soft gradient coupling. This new method, which provides a form of recurrence regularization, can interpolate smoothly between an ordinary ResNet and a "recurrent" ResNet (i.e., one that uses identical weights across layers and thus could be physically implemented with a recurrent network computing the successive stages iteratively across time). We define three indices of recurrent iterative computation and show that a higher degree of gradient coupling promotes iterative convergent computation in ResNets. To measure the degree of weight sharing, we quantify the effective number of parameters of models along the continuum between nonrecurrent and recurrent. For a given effective number of parameters, recurrence regularization does not improve classification accuracy on three visual recognition tasks (MNIST, CIFAR-10, Digitclutter). ResNets, thus, may not benefit from a more similar sets of weights across layers, suggesting that their power does not derive from implicitly recurrent computation.

1 INTRODUCTION

An iterative method solves a difficult estimation or optimization problem by starting from an initial guess and repeatedly applying a transformation that is known to improve the estimate, leading to a sequence of estimates that converges to the solution. Iterative methods provide a powerful approach to finding exact or approximate solutions where direct methods fail (e.g., for difficult inverse problems or solutions to systems of equations that are nonlinear and/or large).

Recurrent neural networks (RNNs) iteratively apply the same transformation to their internal representation, suggesting that they may learn algorithms similar to the iterative methods used in mathematics and engineering. The idea of iterative refinement of a representation has also driven recent progress in the context of feedforward networks. New architectures, based on the idea of iterative refinement, have allowed for the training of very deep feedforward models with hundreds of layers. Prominent architectures for achieving high depth are residual (ResNets; He et al., 2016a) and highway networks (Srivastava et al., 2015), which use skip connections to drive the network to learn the residual, thus encouraging the model to learn successive refinements of a representation of the input that is shared across layers.

One motivation for skip connections is that they alleviate the problem of vanishing or exploding gradients (Hochreiter, 1991). However, the successive refinement of a shared representational format also relates residual and highway networks to RNNs and iterative methods. Learning a single transformation that can be iteratively applied to converge on a good solution promises to provide a useful inductive bias.

However, it is unclear whether recurrent and non-recurrent ResNets indeed learn solutions akin to iterative methods. The two defining features of iterative methods are (1) iteration and (2) convergence. Here we analyze to what extent neural network models learn algorithms that approximate these features. Such iteratively convergent behavior (henceforth simply iterative behavior) may provide a useful inductive bias. We therefore aim to test whether making a network more iteratively convergent can improve its performance (Hypothesis 1).

Recurrence implies iteration, but not convergence, and so is not sufficient for a network to implement an iterative method as defined above. However, since iterative algorithms can in general be cast as recurrent, we expect an association between the two concepts. This suggests that ResNets may employ implicitly recurrent computation (Hypothesis 2). We define a network as implicitly recurrent if its performance can be recovered by a network with similar structure and, in part, explicitly recurrent computations. Such a partly recurrent network would therefore be able to retain a ResNet's performance with fewer parameters.

To test both hypotheses, we manipulate the degree of weight sharing in ResNets, smoothly interpolating between ordinary and recurrent ResNets. Addressing Hypothesis 1, we find that a higher degree of weight sharing tends to make the network more iterative, but does not result in improved performance. Encouraging iterative behavior in this way therefore does not improve the inductive bias. Next, we introduce an effective parameter count that takes into account the degree of weight sharing. Softly coupled networks do not actually require fewer parameters. They may, however, be closer to a fully recurrent computation and the measure we introduced allows us to take this into account. Under this measure, we find that more recurrent networks are not more parameter-efficient, which addresses Hypothesis 2.

2 Related work

Prior theoretical work has focused on explaining the success of ResNets (He et al., 2016a) and the more general class of highway networks (Srivastava et al., 2015) by studying the learning dynamics in ResNets (Hochreiter, 1991; Orhan & Pitkow, 2018; Balduzzi et al., 2017) and their interpretation as an ensemble of shallow networks (Veit et al., 2016; Huang et al., 2018), as a discretized dynamical system (E, 2017; Haber & Ruthotto, 2018; E et al., 2019), and as performing iterative refinement.

The iterative refinement hypothesis. Our work builds on Jastrzebski et al. (2018) who argue that the sequential application of the residual blocks in a ResNet iteratively refines a representational estimate. Their work builds on observations that dropping out residual blocks, shuffling their order, or evaluating the last block several times retains reasonable performance (Veit et al., 2016; Greff et al., 2017) and can be used for training (Huang et al., 2016). Another set of methods uses such perturbations to train deep neural networks, using stochastic depth (Huang et al., 2016; Hu et al., 2019; Press et al., 2020). Other methods learn to evaluate a limited number of layers that depends on the input (Graves, 2017; Figurnov et al., 2017) or learn increasingly fine-grained object categories across layers (Zamir et al., 2017). Instead of using perturbations to encourage stability of the trained network, Ciccone et al. (2018) propose a method inspired by dynamical systems theory to guarantee such stability in their model.

Iterative refinement and inverse problems. The iterative refinement hypothesis is particularly important in the context of inverse problems, which are often solved using iterative methods. This is particularly relevant for ResNets trained for perceptual tasks since perception is often conceptualized as an inverse problem (Poggio et al., 1985; Pizlo, 2001). Rao & Ballard (1999) modeled visual cortex using recurrent neural networks that iteratively infer the latent causes of a hierarchical Bayesian model. Though these networks have been applied to complex datasets (Wen et al., 2018), most models either learn the inverse model (Rezende et al., 2014; Kingma & Welling, 2013) or define an analytically invertible forward model (Rezende & Mohamed, 2015; Gomez et al., 2017). A notable exception are invertible ResNets (Behrmann et al., 2019), whose inverse can be computed through a fixed point iteration. Another set of works starts out with a classical iterative algorithm and unfolds this algorithm to create a deep network, approaching a similar problem from the opposite direction (Gregor & LeCun, 2010; Hershey et al., 2014; Wisdom et al., 2016; 2017). Moreover, Nguyen et al. (2019) refine CNNs to yield an inference algorithm on a specific generative model (recently implemented by Huang et al., 2020).

Recurrent and residual neural networks. The idea of iterative refinement has motivated an increasing number of recurrent convolutional neural networks being applied to computer vision without an explicit implementation of iterative inference (Kubilius et al., 2019; Spoerer et al., 2019). ResNets may be seen as RNNs that have been unfolded for a fixed number of iterations. Sharing weights between the different blocks allows us to train a recurrent residual neural network (Liao & Poggio, 2016). In this framework, recurrent residual neural networks may be seen as a special case of residual neural networks where the weights are equal between all blocks. Savarese & Maire

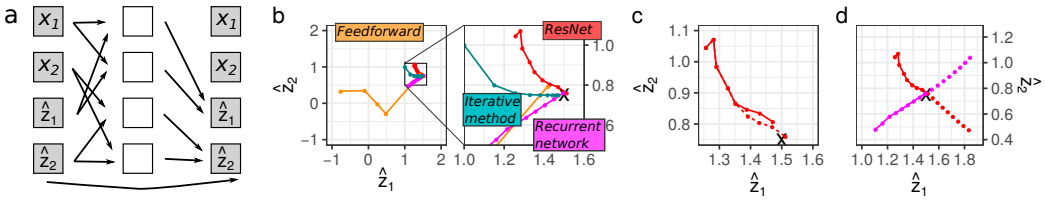

Figure 1: **a** A recurrent ResNet implementing a simple error-corrective inverse model. The prediction based on the current estimate $\hat{\boldsymbol{z}}$ is compared to the input $\boldsymbol{x}$ (via positive and negative errors $\boldsymbol{p}$ and $\boldsymbol{n}$). The error is used to update $\hat{\boldsymbol{z}}$, whereas $\boldsymbol{x}$ is simply retained throughout all blocks. **b** Trajectories of estimates across blocks in a feedforward network, iterative steps in an the error-corrective algorithm, residual blocks in a trained ResNet, and residual blocks in a trained recurrent ResNet. All four methods converge to the correct estimate (indicated by black 'x'). **c** Dropping out the fourth block has a minor impact on the ResNet. **d** If the last block is iteratively applied to the final estimate, the value diverges again, indicating that both the residual and the recurrent ResNet do not learn a convergent structure.

(2019) and Oh et al. (2019) relaxed this constraint by defining the weights of the different blocks as a linear combination of a smaller set of underlying weights. The work by Savarese & Maire (2019) is particularly interesting in this context, as their method yielded a more efficient parameterization of Wide ResNets (Zagoruyko & Komodakis, 2017), which generally have more channels, but far fewer layers than the architectures we consider here.

3 Do ResNets implement iterative computations?

Iterative methods are characterized by two key properties: iteration (i.e., recurrent computation) and convergence. In this section, we examine whether residual neural networks have these properties. We start by showing that ResNets *can express iterative algorithms*, but then demonstrate that ResNets *do not automatically learn iterative algorithms*. We show that this observation extends to large-scale neural networks trained on real data. Finally, we introduce a paradigm that allows us to compare neural networks to iterative methods.

3.1 ResNets can represent iterative computations

A popular application for iterative methods is given by inverse problems of the form $\boldsymbol{x} = f(\boldsymbol{z})$. It is often not possible to directly compute the inverse, f^{-1}, of the forward model, because of analytical or combinatorial intractability. Instead, approximations (Gershman & Goodman, 2014) or iterative error-corrective algorithms (Donoho, 2006) are used. Consider the linear forward model $\boldsymbol{x} = f(\boldsymbol{z}) := (\alpha_1 z_1, \alpha_2 z_2)$. Though f has an analytical inverse, it serves as an illustrative example. Based on an input $\boldsymbol{x}$, we can infer the latent variable $\boldsymbol{z}$ using the iterative update

$$\hat{\boldsymbol{z}}^{(0)} := \boldsymbol{x}, \quad \hat{z}_i^{(t+1)} := \hat{z}_i^{(t)} - s_i \cdot \epsilon_i \cdot (\alpha_i \hat{z}_i^{(t)} - x_i), \quad s_i := \operatorname{sign}(\alpha_i),$$

where $\hat{\boldsymbol{z}}^{(T)}$, for some T, is the estimate of $\boldsymbol{z}$ and $\epsilon_i > 0$ should be sufficiently small[1].

This update can be implemented in a small ResNet. Figure 1a contains a schematic illustration of one block of this network. The representation spanned by an encoder in the beginning of the network contains an input representation $\boldsymbol{x}$ and a representation of the current estimate $\boldsymbol{z}$. In the hidden layer of the residual block, we determine the positive and negative prediction errors and use them to update $\boldsymbol{z}$:

$$p_i^{(t)} := \operatorname{ReLU}(\alpha_i z_i^{(t)} - x_i), \quad n_i^{(t)} := \operatorname{ReLU}(x_i - \alpha_i z_i^{(t)}), \quad z_i^{(t+1)} := z_i^{(t)} - s_i \epsilon_i p_i^{(t)} + s_i \epsilon_i n_i^{(t)}.$$

This recurrent residual building block would implement an iterative computation in a ResNet. The linear model is, of course, a trivial case, but serves as an illustration of the appeal of this approach. A wide neural network is a flexible function approximator and learning to represent the prediction error instead of the prediction is easier in many cases (Zamir et al., 2017).

[1] In the case of a linear function, any $\epsilon_i \leq 1$ works. The example given in figure 1 uses $\epsilon_1 = 0.3$, $\epsilon_2 = 0.8$.

3.2 RESNETS DO NOT AUTOMATICALLY LEARN ITERATIVE COMPUTATIONS

The fact that ResNets can express iterative computations does not imply that they necessarily learn iterative solutions when trained via backpropagation.

Here we consider the simple example from above to better understand the behavior that may emerge. We highlight three behaviors that distinguish iterative from non-iterative computations and will examine the behavior of large-scale neural networks in the following sections. As an example, we train a conventional ResNet, and a ResNet that uses the same weights in each block (equivalent to a recurrent network) to invert the function $f(\boldsymbol{z}) := (\frac{3}{2}z_1, \frac{3}{4}z_2)$. We contrast their behavior with the iterative error-corrective algorithm outlined above.

Due to the lack of constraints, a non-residual feedforward neural network changes its representation in every layer. As a consequence, the linear decoder at the end of the network is not aligned to the intermediate representation and early readout leads to a meaningless estimate. In contrast, the skip connections encourage a ResNet to use the same representational format across blocks. Its intermediate representations are therefore better aligned with the final decoder. Early readout is therefore possible and the representation across blocks will approach the final estimate. As a consequence, the across-block dynamics of the non-residual network are meaningless, whereas the recurrent and residual network's early readouts are close to the final estimate and approach it in a smooth manner, just like the error-corrective algorithm (Figure 1b). Since iterative methods iteratively refine their initial estimate, their behavior is more similar to the ResNets' monotonic convergence.

Besides their smooth convergence, a fundamental property of iterative methods is their recurrence (i.e., the repeated use of the same computation). This means that dropping out an earlier block has the same effect as dropping out the last one. We can relax the requirement for exact repetition (weight sharing) and require merely similar computations. Figure 1c illustrates that the trained ResNet is indeed relatively robust to block dropout.

Yet the learned models fall short of an iterative method, which is apparent from a third mode of investigation. Iterative convergence would imply that applying the last block's transformation iteratively should keep the readout in the vicinity of the actual estimate. This is clearly not the case in our toy example (Figure 1d). Rather than representing a convergent estimate, this result is more compatible with understanding ResNets as approaching their final estimate at a constant speed and, in the case of late readout, moving past this estimate and overshooting at the same speed. Notably, both the non-recurrent and recurrent networks exhibit this behavior.

3.3 ITERATIVE CONVERGENCE INDICES

This behavior is not surprising. After all, the network is trained to work for a fixed number of steps and not constrained to stay within the vicinity of its final estimate if more steps are added. However, it reveals that even recurrent ResNets do not automatically learn iterative convergent computations. To assess the extent to which they do learn iterative convergent computations we define three continuous indices, which measure convergence, recurrence, and divergence (defined below).

We evaluate the indices for six instances of ResNet-104, trained on CIFAR-10 (Krizhevsky, 2009). This ResNet consists of three stages of 16 residual blocks with 16, 32, and 64 channels, respectively.[2] The ResNet achieved 5.2% classification error on the validation set. To characterize the extent to which the ResNets have learned an iterative convergent computation, we introduce three indices measuring different aspects of such computations.

Convergence Index. Viewing ResNets as performing iterative refinement suggests that each stage gradually approaches its final estimate before passing this estimate to the next stage. By passing the estimate at each of the residual blocks to the next stage, we can monitor how the stages approach their final estimate across blocks (see Figure 2a, left panel). In accordance with previous results (Jastrzebski et al., 2018), we find that all stages smoothly approach their final estimate. To measure the rate of convergence, we compute the area under the curve (AUC) of the classification error, which we call the Convergence Index. We invert and normalize this value such that a Convergence

[2]The ResNet-104 uses the architecture recommended by He et al. (2016b). Appendix A includes details on the architecture and the training paradigm and Appendix C corresponding results on two further datasets, MNIST (LeCun et al., 2010) and Digitclutter-3 (Spoerer et al., 2017).

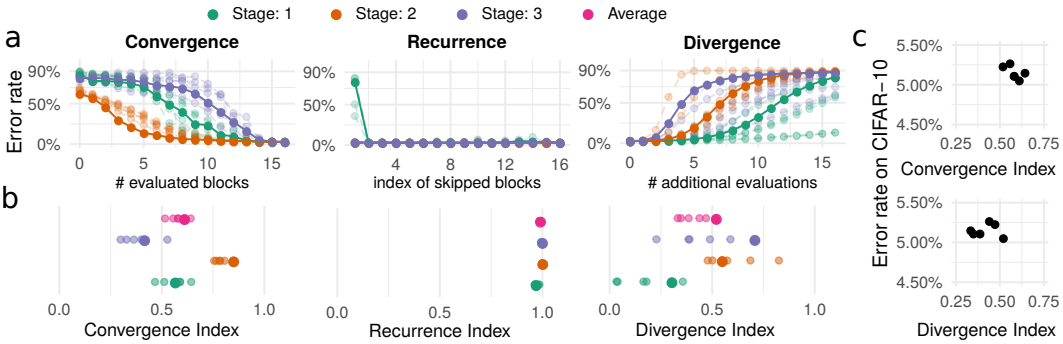

Figure 2: Iterative convergence in ResNets with standard training. **a** The different perturbation methods (early read-out for determining convergence, dropping-out blocks for determining recurrence, and additional evaluations of the last block for determining divergence) are illustrated for the three stages of the ResNet. The x axis depicts the residual block targeted by the perturbation and the y axis the error rate resulting from the corresponding perturbation (chance performance at 90%). For clarity, one of the six instances is emphasized in the plots. **b** The resulting index values for each stage (small translucent dots) and their averages across instances (large dots). **c** The error rate for the individual network instances is plotted against Convergence and Divergence Index.

Index of 0 corresponds chance level read-out at each residual block, whereas a Convergence Index of 1 corresponds to an instant convergence to the final classification error at the first residual block. Figure 2b, left panel, depicts this value for each stage, and averaged across stages.

Recurrence Index. To measure the degree of recurrence, we evaluated the effect of dropping out individual blocks on the error rate of the network (see Figure 2a, middle panel). In a non-recurrent ResNet, dropping out earlier blocks may have a stronger effect on the error rate than dropping out the last block. In contrast, in a recurrent ResNet, the effect on error rate is the same for dropping out either earlier blocks or the last block. We therefore computed the difference in error rate observed after these two manipulations. We summarized the behavior by the AUC, which we refer to as Recurrence Index (RI). We invert and normalize this value such that the RI is 0 if dropping out any block leads to an error rate at chance level and the RI is 1 in the case of a recurrent algorithm. Even though we study non-recurrent ResNets, dropping out any block other than the very first leads to a negligible drop in performance, replicating previous results (Jastrzebski et al., 2018; Veit et al., 2016). As a consequence, the RI, across all stages, is close to 1.0 (see Figure 2b, middle panel).

Divergence Index. ResNets may either converge to their final estimate or simply approach it in a sequence of steps. An iterative algorithm should not be negatively affected by additional applications of the same function. To examine this property, we apply the last block of each stage for an additional up to sixteen steps (see Figure 2a, right panel) and determine the AUC (Divergence Index, DI, see Figure 2b, right panel). We find that no stage is particularly robust to such additional evaluations, though the first stage clearly has the lowest DI, indicating that it is the most robust. This suggests that ResNets approach and move away from their final estimate in a sequence of steps, with their computations bearing little similarity to an iterative convergent algorithm. A high DI does not indicate that the ResNet has failed in some way. After all, it was not trained to be robust to such perturbations. However, it calls into question whether a ResNet implements an iterative convergent computation.

4 Recurrence-regularized residual networks

4.1 Reasons for recurrence regularization

We provided indices measuring convergence, recurrence, and divergence to assess the degree to which a ResNet implements an iterative method. Even though they are able to, ResNets do not necessarily learn to implement a purely iterative method. In particular, they show divergent behavior. This suggests that, for many computational problems, neither a purely feedforward nor a purely

iterative method might be the optimal solution. Instead, some balance between feedforward and iterative computations might have been learned by the trained ResNets.

For our first hypothesis, we examine whether iterative convergence provides a useful inductive bias in ResNets. In principle, this may already be studied using the ResNets discussed above. Specifically, Figure 2c illustrates the relationship between performance and the Convergence and Divergence Index across instances. Due to the lack in variability of the networks' indices and performance, no relationship between iterative convergence and performance is discernible. We therefore introduce soft gradient coupling to manipulate the degree of iterative convergence within ResNets and hypothesize that this may lead to better performance.

We also aim to examine whether these more recurrent networks are more parameter-efficient. This would suggest that part of a ResNet's computations are implicitly recurrent. We therefore examine ResNets with a variable number of parameters. To manipulate parametric complexity without confounding it with computational depth, we vary the number of channels in each stage. We consider networks with 8, 16, 32, or 64 channels in the first stage. The number of channels are doubled in both subsequent stages. We thus cover a wide range of parameter counts, while keeping computational depth constant.

4.2 Soft gradient coupling can interpolate smoothly between ordinary and recurrent optimization

We propose a method to blend between recurrent and non-recurrent networks without changing the architecture or the loss landscape. The method is motivated by the observation that we can train a recurrent neural network by sharing the different blocks' gradients (Jastrzebski et al., 2018; Liao & Poggio, 2016). In ordinary ResNets, the residual block t with the weights $\boldsymbol{W}_t$ is changed by following the gradient $\Delta_t = \partial_{\boldsymbol{W}_t} L$, whereas RNNs impose as gradient

$$\Delta = T^{-1} \sum_{t=1}^{T} \partial_{\boldsymbol{W}_t} L, \tag{1}$$

where the weights across residual blocks within a stage must start from the same initialization. The former means that we do not employ any inductive bias towards recurrence, whereas the latter imposes a possibly overly restrictive function space on the architecture. To address both limitations, we propose **soft gradient coupling**, which uses as its update rule

$$\tilde{\Delta}_t = (1-\lambda)\Delta_t + \lambda\Delta, \quad \lambda \in [0, 1]. \tag{2}$$

For $\lambda < 1$, this retains the entire space of computations enabling both non-recurrent as well as recurrent computations. However, for $\lambda > 0$, the optimization is biased to find more recurrent optima. In contrast to penalty regularizations or strict weight sharing models (Oh et al., 2019; Savarese & Maire, 2019), this does not change the network or loss landscape, but simply the accessibility of different local minima of the loss landscape.

For networks with coupled gradients (i. e., $0 < \lambda \leq 1$), we initialize their weights recurrently (i.e., all residual blocks within one stage share the same initialization). We unshare batch norm statistics as suggested by Jastrzebski et al. (2018), but leave their parameters (softly) coupled.

4.3 Effective parameter count

Softly coupled ResNets have the same number of raw parameters as ordinary ResNets with the same architecture. Their soft weight sharing, however, increases the similarity between weights of different blocks. This weight similarity is not reflected in the raw parameter count. We therefore introduce an **effective parameter count** (EPC). For this measure, three quantities are of interest: the parameters that are not coupled with any other parameters, $\boldsymbol{W}_{\text{uncoupled}}$, the mean parameters for a given stage s, $\overline{\boldsymbol{W}}^{(s)}$, and the normalized deviations from this mean of a parameters within a particular layer, $\boldsymbol{D}_l^{(s)} := (\boldsymbol{W}_l^{(s)} - \overline{\boldsymbol{W}}^{(s)})/\hat{\sigma}(\overline{\boldsymbol{W}}^{(s)})$. As a measure of the effective number of parameters which is sensitive to different degrees of weight sharing, we now propose

$$\text{EPC}(\boldsymbol{W}) := \|\boldsymbol{W}_{\text{uncoupled}}\|_0 + \sum_{s=1}^{S} \|\overline{\boldsymbol{W}}^{(s)}\|_0 + \sum_{s=1}^{S} \sum_{l=1}^{L} \|\boldsymbol{D}_l^{(s)}\|_1, \tag{3}$$

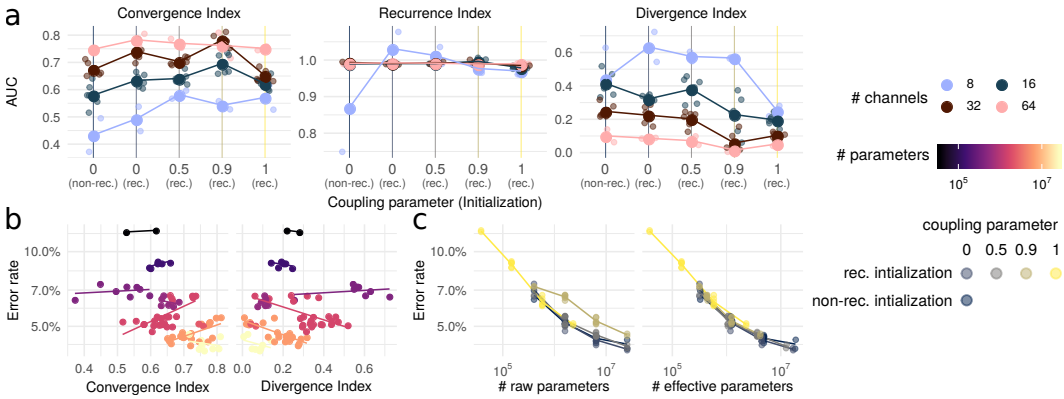

Figure 3: Iterative convergence and performance of coupled ResNets. **a** Effect of gradient coupling and initialization (rec.: recurrent; non.-rec.: non-recurrent) on indices of iterative convergence for architectures with different numbers of channels. **b** Relationship between network performance and measures of iterative convergence, i.e., Convergence (left) and Divergence Index (right). Models with the same number of parameters are visualized by the same color and individual lines. **c** Relationship between the number of raw parameters (left) and the effective parameter count (right) with the error rate for different coupling parameters.

where the L_0-norm $\|\cdot\|_0$ corresponds to the parameter count, and S and L to the number of stages and blocks, respectively. This allows us to interpolate between the raw parameter counts of non-recurrent and recurrent ResNet. For softly coupled ResNets, the EPC measure lies between the raw parameter count of non-recurrent and recurrent ResNets as a function of the weight similarity across weights (see Appendix B.5).

5 Experiments

To assess our hypotheses, we considered non-recurrently initialized (ordinary) ResNets as well as recurrently initialized ResNets with coupling parameters 0, 0.5, 0.9, and 1. We trained several instances of all these ResNets with 8, 16, 32, and 64 channels in the first stage. As before, we here show the results for ResNets trained on CIFAR-10 (Krizhevsky, 2009). We found similar results for MNIST (LeCun et al., 2010) and Digitclutter-3 (Spoerer et al., 2017) (see Appendix C).

5.1 Soft gradient coupling improves Iterative Recurrence Indices

As Figure 3a shows, soft gradient coupling indeed improves iterative convergence, increasing Convergence and Recurrence Indices and decreasing the Divergence Index. Except for a few outliers for the very thin network using 8 channels, the Recurrence Index is centered closely around 1. Convergence and Divergence Index, on the other hand, tend to increase and decrease, respectively, both with higher coupling parameters and with a higher number of channels. Notably, this trend appears to not hold up for a fully recurrent ResNet, corresponding to a coupling parameter of 1. These results show that soft gradient coupling is an effective way of manipulating a ResNet's behavior. Increasing weight similarity across blocks leads to more iterative convergence in a ResNet. Being able to manipulate indices of iterative convergence, we can turn to test our two hypotheses.

5.2 Stronger iterative convergence does not provide a useful inductive bias

We first assessed the effect of gradient coupling on performance. As Table 1 shows, a higher coupling parameter consistently leads to a higher error rate. However, for intermediate coupling parameters of 0.5 and 0.9, this increase in error rate is smaller for networks with higher capacity (more channels). To test the hypothesis whether iterative convergence is a useful inductive bias for ResNets, we assessed the relationship between measures of iterative convergence and performance. Figure 3b demonstrates that performance is higher for ResNets with a higher Convergence

Index and a lower Divergence Index. This effect, however, is driven by the fact that ResNets with a higher number of channels also show higher measures of iterative convergence (see Figure 3a). When controlling for the number of parameters (see individual lines in Figure 3b), we find no relationship between Convergence and Divergence Index and performance. This suggests that iterative convergence is not a useful inductive bias that improves performance in ResNets.

# channels	0 (non-rec.)	0 (rec.)	0.5	0.9	1
8	6.68%	7.01%	6.88%	7.24%	12.1%
16	5.16%	5.19%	5.27%	6.57%	8.9%
32	4.47%	4.41%	4.62%	5.25%	6.28%
64	4.24%	4.04%	4.15%	4.62%	5.09%

Table 1: Mean error rate on the validation dataset of CIFAR-10 across the trained instances for a given coupling parameter and number of channels. Columns denote the coupling parameter of the soft gradient coupling. Note, that $\lambda = 0$ (non-rec.) corresponds to a ResNet with standard training and randomly initialized weights. All other models were trained with recurrent initialization (rec.).

5.3 Stronger gradient coupling does not increase parameter efficiency

We hypothesized that ResNets may employ implicitly recurrent computations. This would mean that there is a more recurrent network which is more parameter-efficient. If this is the case, we would expect soft gradient coupling to find such a solution. For a given number of parameters, however, recurrent and softly gradient-coupled networks do not outperform ordinary ResNets (see Figure 3c, left). Softly coupled networks even tend to do worse. This is because pure parameter counts do not take into account the soft weight sharing across blocks. We therefore developed a measure of the effective parameter count which takes the similarity of weights across layers into account (see Section 4.3). For a given effective parameter count, ordinary, softly gradient-coupled, and recurrent ResNets all perform equally well. Nonetheless, this suggests that we cannot increase parameter efficiency by making the ResNets more recurrent.

6 Discussion

We introduced soft gradient coupling, a new method of recurrence regularization, and demonstrated that this enabled us to manipulate iterative convergence properties in ResNets. To measure these properties, we introduced three indices of iterative convergence, quantifying the effect of perturbations previously introduced in the literature (Veit et al., 2016; Jastrzebski et al., 2018).

Iterative methods are considered powerful approaches in particular for solving difficult inverse problems. However, here we did not find iterative convergence to be a useful inductive bias for ResNets. Moreover, we found that higher degrees of weight sharing did not improve a ResNet's parameter efficiency. One reason for this may be that soft gradient coupling is the wrong method for this purpose or requires a different optimization strategy.[3] Our findings also suggest, however, that deep feedforward computations may not be characterized as iterative refinement on a latent representation, but, at most, as non-iterative refinement on this representation. Our conclusions are based on experiments on three visual classification tasks. Visual tasks are generally considered to be inverse problems and to therefore lend themselves to iterative inference algorithms. However, an iterative method like an error-corrective method that would require a forward model of the data may be more complex and therefore harder to learn than a purely discriminative model. Hence, for the three tested tasks, ResNets may learn direct inference rather than error-corrective inference via a forward model.

Although it did not improve performance here, soft gradient coupling provides a method for smoothly interpolating between feedforward and recurrent neural networks. More generally, soft gradient coupling provides a simple way to encourage sets of weights to remain similar to each other. This technique may find further use in relaxing weight-sharing constraints and studying the benefit of various forms of weight sharing, including recurrence and convolution, in deep neural networks.

[3] Appendix B contains a few variations on the techniques introduced above that do not change our results.

REFERENCES

David Balduzzi, Marcus Frean, Lennox Leary, J. P. Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 342–350. JMLR. org, 2017.

Jens Behrmann, Will Grathwohl, Ricky T. Q. Chen, David Duvenaud, and Jörn-Henrik Jacobsen. Invertible Residual Networks. In *ICML*, pp. 573–582, 2019. URL http://proceedings.mlr.press/v97/behrmann19a.html.

Marco Ciccone, Marco Gallieri, Jonathan Masci, Christian Osendorfer, and Faustino Gomez. NAIS-Net: Stable Deep Networks from Non-Autonomous Differential Equations. *arXiv:1804.07209 [cs, stat]*, November 2018. URL http://arxiv.org/abs/1804.07209. arXiv: 1804.07209.

Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, and Aaron Courville. Recurrent batch normalization. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006. Publisher: IEEE.

Weinan E. A Proposal on Machine Learning via Dynamical Systems. *Communications in Mathematics and Statistics*, 5(1):1–11, March 2017. ISSN 2194-671X. doi: 10.1007/s40304-017-0103-z. URL https://doi.org/10.1007/s40304-017-0103-z.

Weinan E, Chao Ma, and Lei Wu. Barron Spaces and the Compositional Function Spaces for Neural Network Models. *arXiv:1906.08039 [cs, math, stat]*, June 2019. URL http://arxiv.org/abs/1906.08039. arXiv: 1906.08039.

Michael Figurnov, Maxwell D. Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially Adaptive Computation Time for Residual Networks. pp. 1039–1048, 2017. URL http://openaccess.thecvf.com/content_cvpr_2017/html/Figurnov_Spatially_Adaptive_Computation_CVPR_2017_paper.html.

Samuel Gershman and Noah Goodman. Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society*, volume 36, 2014. Issue: 36.

Aidan N. Gomez, Mengye Ren, Raquel Urtasun, and Roger B. Grosse. The Reversible Residual Network: Backpropagation Without Storing Activations. *arXiv:1707.04585 [cs]*, July 2017. URL http://arxiv.org/abs/1707.04585. arXiv: 1707.04585.

Alex Graves. Adaptive Computation Time for Recurrent Neural Networks. *arXiv:1603.08983 [cs]*, February 2017. URL http://arxiv.org/abs/1603.08983. arXiv: 1603.08983.

Klaus Greff, Rupesh K. Srivastava, and Jürgen Schmidhuber. Highway and Residual Networks learn Unrolled Iterative Estimation. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pp. 399–406, Madison, WI, USA, June 2010. Omnipress. ISBN 978-1-60558-907-7.

Eldad Haber and Lars Ruthotto. Stable Architectures for Deep Neural Networks. *Inverse Problems*, 34(1):014004, January 2018. ISSN 0266-5611, 1361-6420. doi: 10.1088/1361-6420/aa9a90. URL http://arxiv.org/abs/1705.03341. arXiv: 1705.03341.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. pp. 1026–1034, 2015. URL https://www.cv-foundation.org/openaccess/content_iccv_2015/html/He_Delving_Deep_into_ICCV_2015_paper.html.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016a.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016b.

John R. Hershey, Jonathan Le Roux, and Felix Weninger. Deep Unfolding: Model-Based Inspiration of Novel Deep Architectures. *arXiv:1409.2574 [cs, stat]*, November 2014. URL http://arxiv.org/abs/1409.2574. arXiv: 1409.2574.

Sepp Hochreiter. *Untersuchungen zu dynamischen neuronalen Netzen*. PhD thesis, 1991.

Hanzhang Hu, Debadeepta Dey, Martial Hebert, and J. Andrew Bagnell. Learning Anytime Predictions in Neural Networks via Adaptive Loss Balancing. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):3812–3821, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.33013812. URL https://www.aaai.org/ojs/index.php/AAAI/article/view/4268. Number: 01.

Furong Huang, Jordan Ash, John Langford, and Robert Schapire. Learning Deep ResNet Blocks Sequentially using Boosting Theory. In *International Conference on Machine Learning*, pp. 2058–2067, 2018.

Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep Networks with Stochastic Depth. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *Computer Vision – ECCV 2016*, Lecture Notes in Computer Science, pp. 646–661, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46493-0. doi: 10.1007/978-3-319-46493-0_39.

Yujia Huang, James Gornet, Sihui Dai, Zhiding Yu, Tan Nguyen, Doris Y. Tsao, and Anima Anandkumar. Neural Networks with Recurrent Generative Feedback. *arXiv:2007.09200 [cs, stat]*, July 2020. URL http://arxiv.org/abs/2007.09200. arXiv: 2007.09200.

Yerlan Idelbayev. Proper ResNet Implementation for CIFAR10/CIFAR100 in PyTorch, 2018. URL https://github.com/akamaster/pytorch_resnet_cifar10.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

Stanisław Jastrzebski, Devansh Arpit, Nicolas Ballas, Vikas Verma, Tong Che, and Yoshua Bengio. Residual Connections Encourage Iterative Inference. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=SJa9iHgAZ.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. Technical report, 2009.

Jonas Kubilius, Martin Schrimpf, Kohitij Kar, Rishi Rajalingham, Ha Hong, Najib Majaj, Elias Issa, Pouya Bashivan, Jonathan Prescott-Roy, Kailyn Schmidt, Aran Nayebi, Daniel Bear, Daniel L Yamins, and James J DiCarlo. Brain-Like Object Recognition with High-Performing Shallow Recurrent ANNs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 12805–12816. Curran Associates, Inc., 2019. URL http://papers.nips.cc/paper/9441-brain-like-object-recognition-with-high-performing-shallow-recurrent-anns.pdf.

Yann LeCun, Corinna Cortes, and C. J. Burges. MNIST handwritten digit database. 2010.

Qianli Liao and Tomaso Poggio. Bridging the gaps between residual learning, recurrent neural networks and visual cortex. *arXiv preprint arXiv:1604.03640*, 2016.

Tan Nguyen, Nhat Ho, Ankit Patel, Anima Anandkumar, Michael I. Jordan, and Richard G. Baraniuk. A Bayesian Perspective of Convolutional Neural Networks through a Deconvolutional Generative Model. *arXiv:1811.02657 [cs, stat]*, December 2019. URL http://arxiv.org/abs/1811.02657. arXiv: 1811.02657.

Jeeheh Oh, Jiaxuan Wang, Shengpu Tang, Michael W. Sjoding, and Jenna Wiens. Relaxed Parameter Sharing: Effectively Modeling Time-Varying Relationships in Clinical Time-Series. In *Machine Learning for Healthcare Conference*, pp. 27–52, October 2019. URL http://proceedings.mlr.press/v106/oh19a.html. ISSN: 1938-7228 Section: Machine Learning.

Travis E Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.

Emin Orhan and Xaq Pitkow. Skip Connections Eliminate Singularities. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=HkwBEMWCZ.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

Thomas Lin Pedersen. *patchwork: The Composer of Plots*. 2019. URL https://CRAN.R-project.org/package=patchwork.

Zygmunt Pizlo. Perception viewed as an inverse problem. *Vision Research*, 41(24):3145–3161, November 2001. ISSN 0042-6989. doi: 10.1016/S0042-6989(01)00173-0. URL http://www.sciencedirect.com/science/article/pii/S0042698901001730.

Tomaso Poggio, Vincent Torre, and Christof Koch. Computational vision and regularization theory. *Nature*, 317(6035):314–319, September 1985. ISSN 1476-4687. doi: 10.1038/317314a0. URL https://www.nature.com/articles/317314a0. Number: 6035 Publisher: Nature Publishing Group.

Ofir Press, Noah A. Smith, and Omer Levy. Improving Transformer Models by Reordering their Sublayers. *arXiv:1911.03864 [cs]*, April 2020. URL http://arxiv.org/abs/1911.03864. arXiv: 1911.03864.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2019. URL https://www.R-project.org/.

Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, January 1999. ISSN 1546-1726. doi: 10.1038/4580. URL https://www.nature.com/articles/nn0199_79.

Danilo Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. In *International Conference on Machine Learning*, pp. 1530–1538, 2015.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning*, pp. 1278–1286, 2014.

Pedro Savarese and Michael Maire. Learning Implicitly Recurrent CNNs Through Parameter Sharing. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=rJgYxn09Fm.

Andri Signorell. *DescTools: Tools for Descriptive Statistics*. 2020. URL https://cran.r-project.org/package=DescTools.

Courtney J. Spoerer, Patrick McClure, and Nikolaus Kriegeskorte. Recurrent Convolutional Neural Networks: A Better Model of Biological Object Recognition. *Frontiers in Psychology*, 8, 2017. Publisher: Frontiers.

Courtney J. Spoerer, Tim C. Kietzmann, and Nikolaus Kriegeskorte. Recurrent networks can recycle neural resources to flexibly trade speed for accuracy in visual recognition. *bioRxiv*, pp. 677237, June 2019. doi: 10.1101/677237. URL https://www.biorxiv.org/content/10.1101/677237v2. Publisher: Cold Spring Harbor Laboratory Section: New Results.

Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.

Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009. ISBN 1-4414-1269-7.

Andreas Veit, Michael J Wilber, and Serge Belongie. Residual Networks Behave Like Ensembles of Relatively Shallow Networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 550–558. Curran Associates, Inc., 2016. URL http://papers.nips.cc/paper/6556-residual-networks-behave-like-ensembles-of-relatively-shallow-networks.pdf.

Haiguang Wen, Kuan Han, Junxing Shi, Yizhen Zhang, Eugenio Culurciello, and Zhongming Liu. Deep Predictive Coding Network for Object Recognition. In *International Conference on Machine Learning*, pp. 5266–5275, 2018.

Hadley Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. ISBN 978-3-319-24277-4. URL https://ggplot2.tidyverse.org.

Hadley Wickham and Lionel Henry. *tidyr: Tidy Messy Data*. 2019. URL https://CRAN.R-project.org/package=tidyr.

Hadley Wickham, Romain François, Lionel Henry, and Kirill Müller. *dplyr: A Grammar of Data Manipulation*. 2019. URL https://CRAN.R-project.org/package=dplyr.

Scott Wisdom, Thomas Powers, James Pitton, and Les Atlas. Interpretable Recurrent Neural Networks Using Sequential Sparse Recovery. *arXiv:1611.07252 [cs, stat]*, November 2016. URL http://arxiv.org/abs/1611.07252. arXiv: 1611.07252.

Scott Wisdom, Thomas Powers, James Pitton, and Les Atlas. Deep recurrent NMF for speech separation by unfolding iterative thresholding. In *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, pp. 254–258, October 2017. doi: 10.1109/WASPAA.2017.8170034. ISSN: 1947-1629.

Sergey Zagoruyko and Nikos Komodakis. Wide Residual Networks. *arXiv:1605.07146 [cs]*, June 2017. URL http://arxiv.org/abs/1605.07146. arXiv: 1605.07146.

Amir R Zamir, Te-Lin Wu, Lin Sun, William B Shen, Bertram E Shi, Jitendra Malik, and Silvio Savarese. Feedback networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1308–1317, 2017.

A Network training

A.1 Software

We trained the networks using PyTorch (Paszke et al., 2019), PytorchLightning, numpy (Oliphant, 2006), and pandas in Python (Van Rossum & Drake, 2009). The remaining analysis was conducted in R (R Core Team, 2019), using ggplot2 (Wickham, 2016), dplyr (Wickham et al., 2019), tidyr (Wickham & Henry, 2019), patchwork (Pedersen, 2019), and DescTools (Signorell, 2020). The implementation of ResNet-104 was significantly supported by an existing implementation (Idelbayev, 2018).

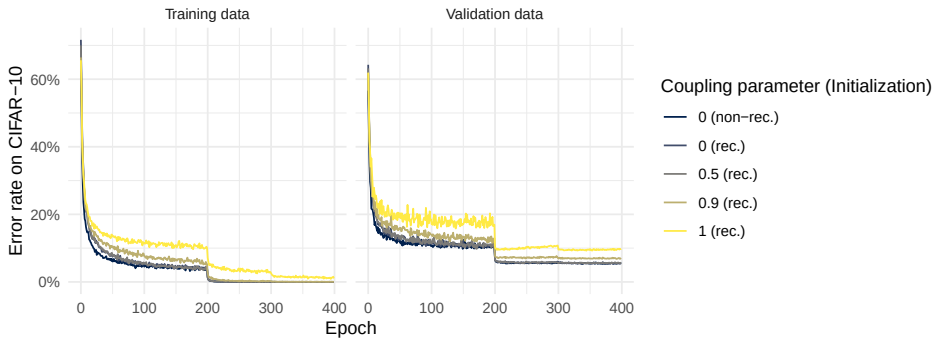

Figure 4: This figure depicts the evolution of mean error rate on training (left) and validation data (right) during training. This pertains to 16 channels in the first stage across all combinations of coupling parameter and initialization.

A.2 Datasets

We trained and evaluated the networks on three datasets: CIFAR-10 (Krizhevsky, 2009), MNIST (LeCun et al., 2010), and Digitclutter (Spoerer et al., 2017). All images were normalized before being provided to the network. For CIFAR-10, we used a training/validation sample split of 45000/5000, augmented the data by random cropping and random horizontal flips during training. For MNIST, we used a split of 50000/10000, and augmented the data by random cropping. Finally, for Digitclutter, we used three overlapping digits (referred to as Digitclutter-3) with a split of 90000/10000 without data augmentation.

A.3 Architecture details

As our base architecture we use a ResNet with the preactivation unit recommended by He et al. (2016b). This means that a residual block consists of

$$\mathrm{BN} \rightarrow \mathrm{ReLU} \rightarrow \mathrm{Conv} \rightarrow \mathrm{BN} \rightarrow \mathrm{ReLU} \rightarrow \mathrm{Conv},$$

where BN stands for batch normalization (Ioffe & Szegedy, 2015), ReLU for rectified linear unit, and Conv for a convolutional layer.

For CIFAR-10 and Digitclutter, this ResNet contained three stages with 16 blocks each. At each stage, width and height of the representation were halved and the number of channels was doubled. This was done with an additional downsampling block. This means that for a standard ResNet with 16 channels in the first stage, the second stage consisted of (16x16)-representations with 32 channels, and the third stage consisted of (8x8)-representations with 64 channels. In contrast, we only used one stage with 16 blocks for MNIST.

A.4 Training details

We initialized our networks using Kaiming initialization (He et al., 2015) for the convolutional and linear weights. In the case of a recurrent initialization, these were equal across residual blocks of a stage. In the case of a non-recurrent initialization, they were drawn independently for each block. We initialized the batch normalization's scale γ with 1 and its bias with 0.

We trained all networks for 400 epochs using gradient descent with momentum 0.9. For CIFAR-10, we used an initial learning rate of 0.1, which was divided by 10 after 200 and 300 epochs (see figure 4 for the evolution of the learning rate of the ResNets with 16 channels in the first stage). For MNIST, we used an initial learning rate of 0.025, which was divided by 2.5 after 200 and 300 epochs. For Digitclutter, we used an initial learning rate of 0.05, which was divided by 5 after 200 and 300 epochs. We then identified the epoch at which the network had obtained the best validation classification error and used the corresponding model for all analyses.

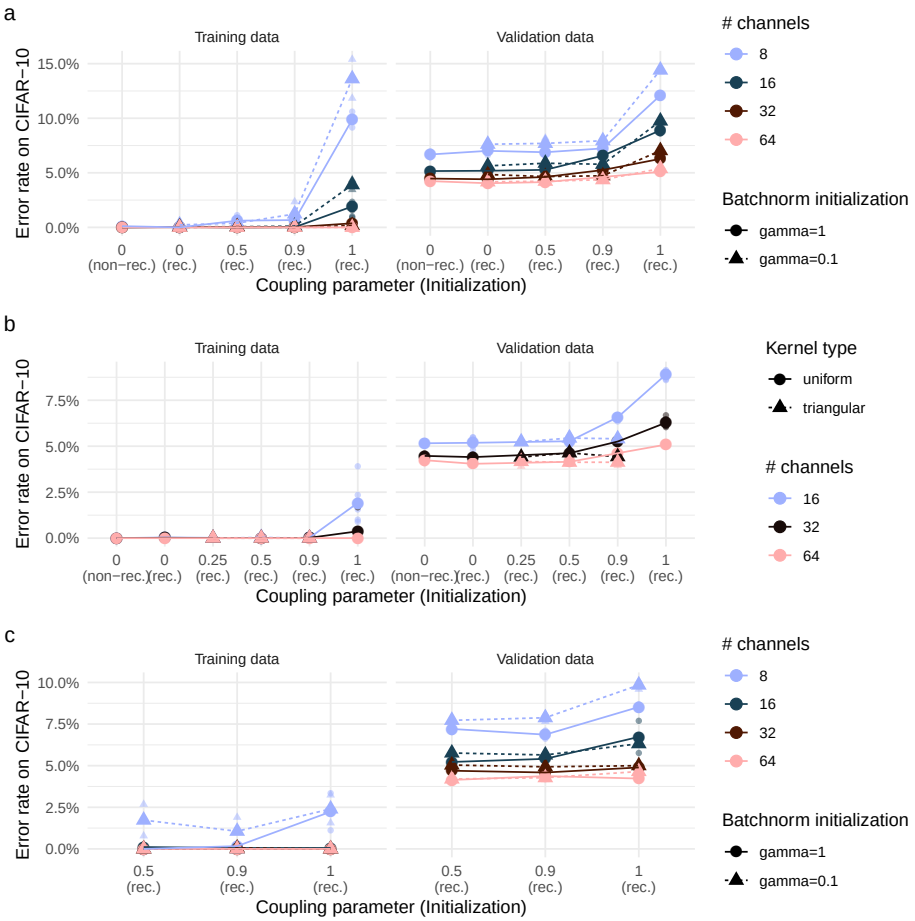

Figure 5: Performance of training variations on CIFAR-10. **a** The effect of initializing batchnorm with $\gamma = 0.1$ instead of $\gamma = 1$. **b** The effect of using a triangular kernel for gradient coupling instead of a uniform kernel. **c** A variation of gradient coupling where the first five blocks in each stage were uncoupled.

B Training variations

On CIFAR-10, we explored several variations of the training algorithm, which were consistent with the findings presented above.

B.1 Recurrent batch normalization

Cooijmans et al. (2017) proposed to initialize the scale of the batch normalization γ as 0.1 in recurrent neural networks. When applying this method to softly gradient-coupled networks, we find that the method strongly improves performance for $\lambda = 0.9$, but not for any other coupling parameters (see figure 5a).

B.2 Triangular gradient coupling

We may generalize soft gradient coupling to the coupling rule

$$\tilde{\Delta}_t = \sum_{s=1}^{T} \kappa(s, t; \lambda) \partial_{\boldsymbol{W}_t} L, \tag{4}$$

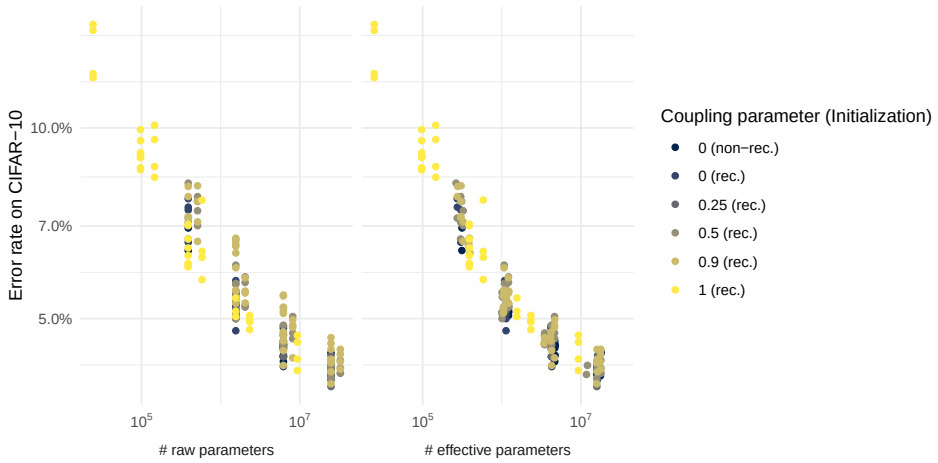

Figure 6: This figure plots the error rate of all models trained on CIFAR-10 against their raw and effective parameter count.

where κ is some kernel depending on the coupling parameter λ. Ordinary soft gradient coupling can be recovered with a uniform kernel

$$\kappa(s, t; \lambda) := \begin{cases} (1-\lambda) + \lambda/T & \text{if } s = t, \\ \lambda/T & \text{if } s \neq t. \end{cases} \tag{5}$$

Alternatively, we may want to couple adjacent blocks more strongly than blocks that are far away from each other. This can, for example, be achieved with a triangular kernel

$$\kappa(s, t; \lambda) := \begin{cases} (1 - 2 \cdot (1-\lambda)/T \cdot |s-t|)^+ & \text{if } \lambda \geq 0.5, \\ (1 - 1/(2 \cdot \lambda \cdot T) \cdot |s-t|)^+ & \text{if } \lambda < 0.5, \end{cases} \tag{6}$$

where

$$(\cdot)^+ = \max(\cdot, 0).$$

Again, $\lambda = 0$ corresponds to an ordinary ResNet and $\lambda = 1$ corresponds to a fully recurrent ResNet, whereas intermediate values for λ smoothly interpolate between the two.

Figure 5b shows the effect of triangular as compared to uniform gradient coupling. For a given coupling parameter, this method appears to outperform the uniformly gradient-coupled network, but neither method outperforms an ordinary ResNet.

B.3 Start coupling at later layers

Finally, we uncoupled the first five residual blocks and only constrained the remaining blocks. For the fully recurrent case, this corresponds to the implementation by Jastrzebski et al. (2018). We also studied the effect of this late coupling together with a batchnorm initialization of $\gamma = 0.1$. In both cases, the late coupling means that the coupled networks are better than coupled networks which start coupling at block 0, but still worse than an ordinary ResNet.

B.4 Parameter efficiency

Figure 6 shows that across all these extended networks, the relationship between effective parameter count and performance is also largely unaffected by coupling parameter.

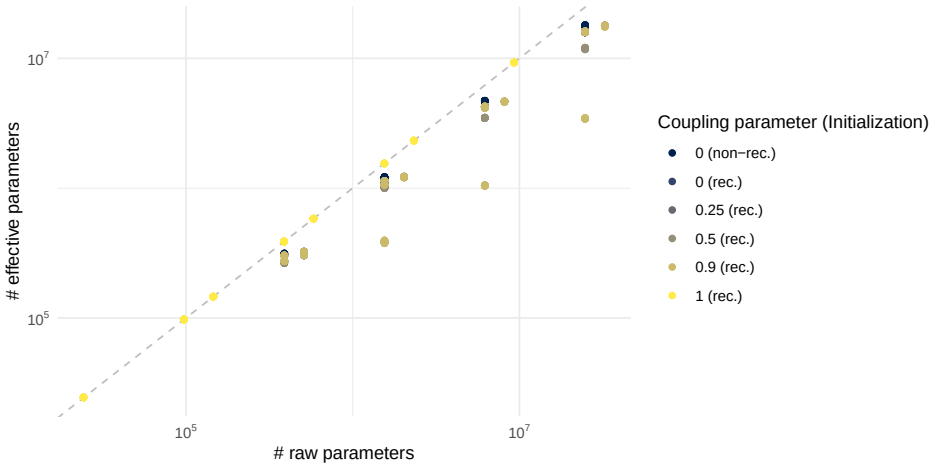

Figure 7: This figure plots the effective parameter count of the networks studied in the main text against the raw parameter count.

B.5 Effective parameter count

Figure 7 plots the effective parameter count of the networks studied in the main text against the raw parameter count. We find that fully recurrent ResNets align perfectly and ordinary, non-recurrent ResNets align almost perfectly, whereas softly-coupled ResNets have a higher raw parameter count than effective parameter count.

B.6 Iterative convergence indices

Figure 8 demonstrates that for the extended training variations, higher coupling parameters generally tend to increase the Convergence Index and decrease the Divergence Index, as well. The extended training variations also predominantly yield networks with a Recurrence Index around 1.

C Results on MNIST and Digitclutter

Tables 2 and 3 demonstrate that, just as for CIFAR-10, increased coupling parameters do not improve performance on MNIST or Digitclutter-3. In contrast to CIFAR-10, we observe that a ResNet with 64 channels obtains worse performance than a ResNet with 32 channels on MNIST. This effect does not hold up for softly coupled networks.

# channels	0 (non-rec.)	0 (rec.)	0.5 (rec.)	0.9 (rec.)	1 (rec.)
16	0.336 %	0.346 %	0.45 %	0.475 %	0.593 %
32	0.277 %	0.341 %	0.361 %	0.401 %	0.386 %
64	0.321 %	0.326 %	0.326 %	0.326 %	0.341 %

Table 2: Results of the trained ResNets on MNIST. Columns refer to the gradient coupling parameter λ and - in brackets - the initialization (non-rec.: non-recurrent initialization; rec.: recurrent initialization

# channels	0 (non-rec.)	0 (rec.)	0.5 (rec.)	0.9 (rec.)	1 (rec.)
16	1.553 %	1.642 %	1.671 %	1.553 %	22.053 %
32	1.068 %	1.216 %	1.276 %	1.291 %	6.779 %
64	1.058 %	1.202 %	1.142 %	1.197 %	2.645 %

Table 3: Results of the trained ResNets on Digitclutter

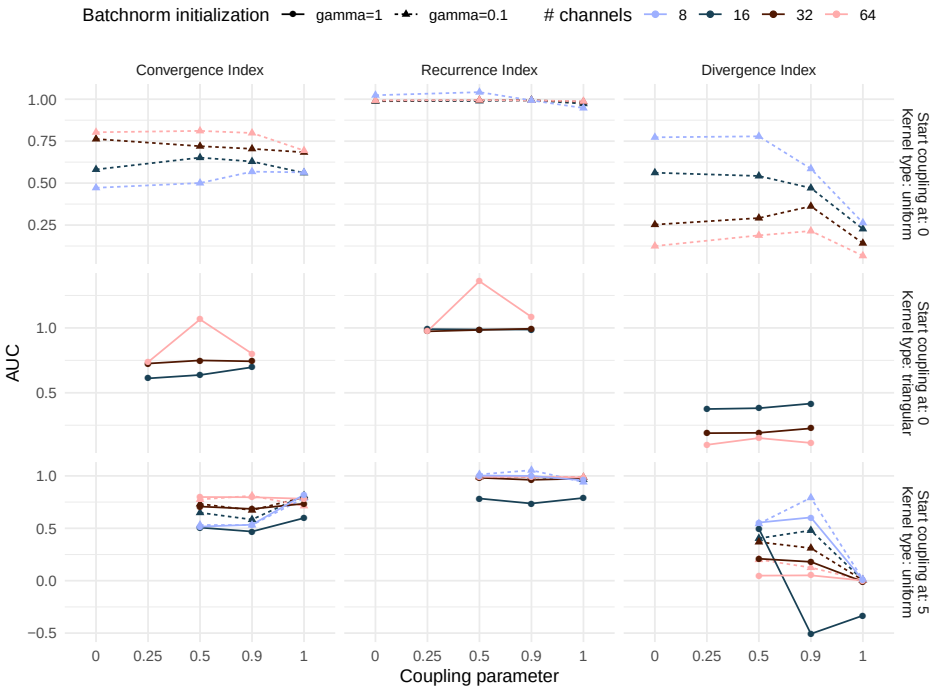

Figure 8: This figure plots the indices of iterative convergence against the coupling parameters for the different training variations.