
Lookaround Optimizer: k steps around, 1 step average

— *Supplementary Material* —

Anonymous Author(s)

Affiliation

Address

email

1 A Exploration in the Loss Landscape

2 To demonstrate the effectiveness of Lookaround during training, we set up an experiment and visualize
3 the loss landscape of the models under different data augmentations in Figure 1. By observing the
4 loss landscape, we can gain a clearer understanding of the role played by the weight averaging at
5 different stages during the training process.

6 **Training Process and Parameter Settings.** We first train a ResNet50 on the CIFAR100 dataset.
7 The learning rate is initialized to 0.1 and decay at 60, 120, and 160 epochs using a MultiStepLR
8 scheduler with a decay factor 0.2. The batch size is set to 128, we use stochastic gradient descent
9 with momentum to optimize the model and use random crops and random vertical flips augmentation
10 to enhance the training datasets. We use model checkpoints at epochs 50, 110, and 150 as our three
11 pretrained models(V-network). The three pretrained models correspond to learning rates of 0.1, 0.02
12 and 0.004, respectively. Using these pretrained models as the starting point, we finetune the model
13 with 1, 10, 100, 1000, 10000 iterations under the corresponding learning rate and the setting of
14 random horizontal flipping (H-network) or RandAugment (R-network).

15 **Visualization Method.** We use the visualization method in [2] to plot the loss landscape. In
16 this method, the weights of the three models are flatten respectively as one-dimensional vectors
17 w_v, w_h, w_r , and then two orthogonal vectors are calculated between the three vectors as the X-axis
18 direction and the Y-axis direction: $u = (w_h - w_v), v = (w_r - w_v, w_h - w_v) / \|w_h - w_v\|^2 \cdot (w_h - w_v)$.
19 Then the normalized vectors $\hat{u} = u / \|u\|, \hat{v} = v / \|v\|$ foam an orthonormal basis in the plane contain
20 w_v, w_h, w_r . Then a point P with coordinates (x, y) in the plane would be given by $P = w_v + x \cdot \hat{u} + y \cdot \hat{v}$.

21 **Discussion and Inspiring.** Under different learning rates and different around steps k , Lookaround
22 has the tendency to lead the model trained on different data augmentation to the near-constant loss
23 manifold. In such circumstances, the "average step" can lead the model into the center of the loss
24 basin to get a lower test loss. However, weight averaging does not necessarily work in all cases. For
25 example, the network obtained after weight averaging gets a larger loss under a large learning rate
26 with a large around step (e.g., $lr = 0.1, k = 10000$). In this case, the model is located on a peak
27 between different basins rather than in the center of a basin. Moreover, in the case of around step
28 $k = 1$, the weight averaging also does not achieve better performance. Nevertheless, such extreme
29 cases do not prevent weight averaging from being a useful tool to speed up the training process. The
30 center of the basin in loss landscape, which requires multi-step gradient descent to reach, can be
31 reached by only one weight averaging step. At other learning rates and around steps k , the models
32 after weight averaging all result in a lower loss than the individual model. Such phenomena encourage
33 us to explore more methods to find more optimal solutions in the loss landscape in the future.

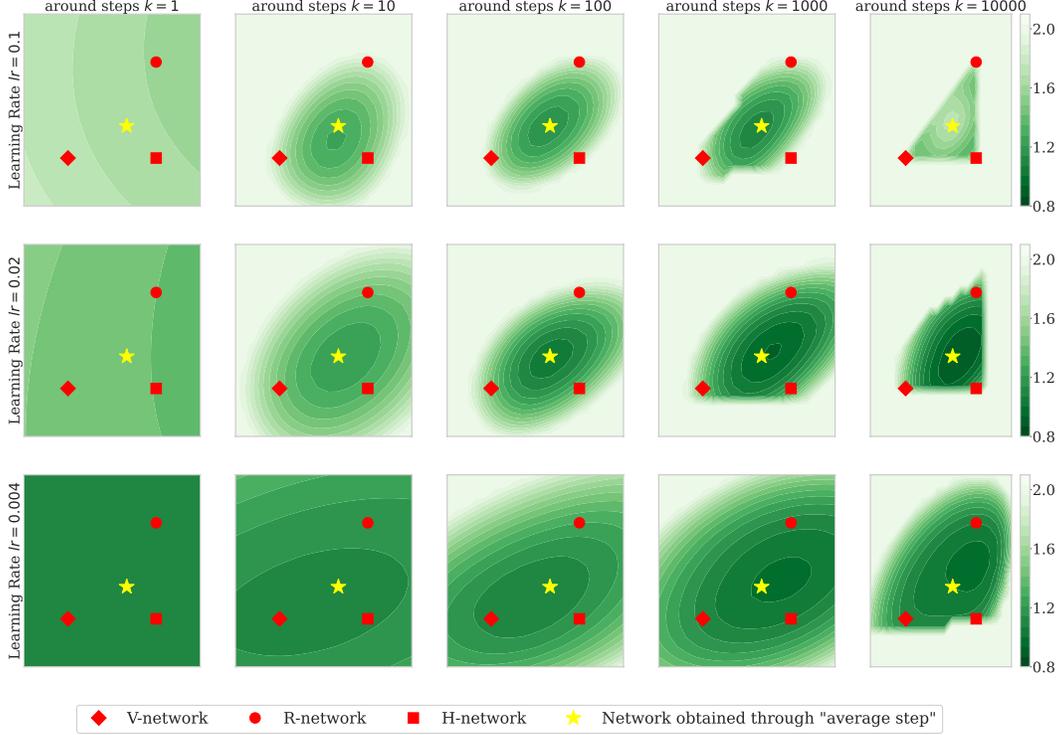


Figure 1: The test loss landscape of ResNet50 on CIFAR100. The diamond block (V-network) represents the pretrained model trained with random vertical flipping. Then we can use random horizontal flipping and RandAugment to finetune V-network to get H-network and R-network.

34 B Steady-state and Convergence Analysis of Lookaround

35 We use quadratic functions to analyze the steady-state and convergence analysis of Lookaround. First,
36 we present the proof of Proposition 1.

37 **Proposition 1** (Steady-state risk). *Let $0 < \gamma < 1/L$ be the learning rate satisfying $L = \max_i a_i$. One
38 can obtain that in the noisy quadratic setup, the variances of SGD, Lookahead [7] and Lookaround
39 converge to the following fixed matrix:*

$$V_{SGD}^* = \frac{\gamma^2 \mathbf{A}^2 \Sigma^2}{\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^2}, \quad (1)$$

$$V_{Lookahead}^* = \underbrace{\frac{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k})}{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k}) + 2\alpha(1 - \alpha)(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^k)}}_{\approx \mathbf{I}, \text{ if } \alpha \in (0,1)} V_{SGD}^*, \quad (2)$$

$$V_{Lookaround}^* = \underbrace{\frac{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k}) + 2\alpha(1 - \alpha)(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^k)}{\alpha^2 (d\mathbf{I} - (d-1)(\mathbf{I} - \gamma \mathbf{A})^{2k})}}_{\approx \mathbf{I}, \text{ if } d \geq 3 \text{ and } \alpha \in [1/2, 1)} V_{Lookahead}^*. \quad (3)$$

40 respectively, where α denotes the average weight factor of models with varying trajectory points.

41 From Wu et al. [6], we have the following conclusions about the dynamics of SGD with learning rate
42 γ :

$$\begin{aligned} \mathbb{E}[\theta^{(t+1)}] &= (\mathbf{I} - \gamma \mathbf{A}) \mathbb{E}[\theta^{(t)}], \\ \mathbb{V}[\theta^{(t+1)}] &= (\mathbf{I} - \gamma \mathbf{A})^2 \mathbb{V}[\theta^{(t)}] + \gamma^2 \mathbf{A}^2 \Sigma. \end{aligned}$$

43 **Lemma 1.** *The expectation and variance of lookaround have the following iterates:*

$$\mathbb{E}[\phi^{(t+1)}] = (\mathbf{I} - \gamma \mathbf{A})^k \mathbb{E}[\phi^{(t)}], \quad (4)$$

$$\mathbb{V}[\phi^{(t+1)}] = \frac{d-1}{d} (\mathbf{I} - \gamma \mathbf{A})^{2k} \mathbb{V}[\phi^{(t)}] + \frac{\gamma^2 \mathbf{A}^2 \Sigma (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k})}{d(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^2)}. \quad (5)$$

44 ϕ_t yields ϕ_{t+1} by performing an around step and an average step.

45 *Proof.* The expected iterate follows from SGD:

$$\begin{aligned} \mathbb{E}[\phi^{t+1}] &= \mathbb{E}\left[\frac{1}{d} \sum_i \theta_{t,i,k}\right] = \sum_i \frac{1}{d} \mathbb{E}[\theta_{t,i,k}] \\ &= \sum_i \frac{1}{d} (\mathbf{I} - \gamma \mathbf{A})^k \mathbb{E}[\theta_{t,i,0}] = (\mathbf{I} - \gamma \mathbf{A})^k \mathbb{E}[\phi^t]. \end{aligned}$$

46 For the variance of ϕ^{t+1} , we can break it down into two parts as follows:

$$\mathbb{V}[\phi^{t+1}] = \mathbb{V}\left[\frac{1}{d} \sum_i \theta_{t,i,k}\right] = \sum_i \frac{1}{d^2} \mathbb{V}[\theta_{t,i,k}] + \sum_{i \neq j, 1 \leq i, j \leq d} \frac{1}{d^2} \text{cov}(\theta_{t,i,k}, \theta_{t,j,k}).$$

47 The covariance of the different models can be calculated in the following way:

$$\begin{aligned} \text{cov}(\theta_{t,i,k}, \theta_{t,j,k}) &= \mathbb{E}[\theta_{t,i,k} \theta_{t,j,k}] - \mathbb{E}[\theta_{t,i,k}] \mathbb{E}[\theta_{t,j,k}] \\ &= \mathbb{E}[(\mathbf{I} - \gamma \mathbf{A})^{2k} (\phi^t)^2] - (\mathbf{I} - \gamma \mathbf{A})^{2k} \mathbb{E}[\phi^t]^2 \\ &= (\mathbf{I} - \gamma \mathbf{A})^{2k} \mathbb{V}[\phi^t]. \end{aligned}$$

48 After permuting and regrouping again, we can obtain the iterate with respect to the variance.

$$\begin{aligned} \mathbb{V}[\phi^{t+1}] &= \sum_{i \neq j, 1 \leq i, j \leq d} \frac{1}{d^2} \text{cov}(\theta_{t,i,k}, \theta_{t,j,k}) + \sum_i \frac{1}{d^2} \mathbb{V}[\theta_{t,i,k}] \\ &= \frac{1}{d^2} (d^2 - d) (\mathbf{I} - \gamma \mathbf{A})^{2k} \mathbb{V}[\phi^t] + \frac{1}{d} \left[\sum_{i=0}^{k-1} (\mathbf{I} - \gamma \mathbf{A})^{2i} \gamma^2 \mathbf{A}^2 \Sigma \right] \\ &= \frac{d-1}{d} (\mathbf{I} - \gamma \mathbf{A})^{2k} \mathbb{V}[\phi^t] + \frac{\gamma^2 \mathbf{A}^2 \Sigma (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k})}{d(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^2)}. \end{aligned}$$

49 The proof is now complete. □

50 **Remark.** From Equation 4, the expectation term for ϕ in Lookaround eventually converges to 0, as
51 does Lookahead and SGD.

52 From Zhang et al. [7], we have the following analysis about the variance fixed point of lookahead
53 with learning rate γ and weight factor α , which represents the average weight factor of models with
54 different trajectory points in the Lookahead optimizer, which is generally (0, 1):

$$V_{\text{Lookahead}}^* = \frac{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k})}{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k}) + 2\alpha(1 - \alpha)(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^k)} V_{\text{SGD}}^*. \quad (6)$$

55 We now derive the fixed point of the variance, proceed with the proof of Proposition 1:

$$V_{\text{Lookaround}}^* = \frac{\alpha^2 (\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^{2k}) + 2\alpha(1 - \alpha)(\mathbf{I} - (\mathbf{I} - \gamma \mathbf{A})^k)}{\alpha^2 (d\mathbf{I} - (d-1)(\mathbf{I} - \gamma \mathbf{A})^{2k})} V_{\text{Lookahead}}^*.$$

Proof.

$$\begin{aligned}
V_{Lookaround}^* &= \frac{d-1}{d}(\mathbf{I} - \gamma\mathbf{A})^{2k}V_{Lookaround}^* + \frac{\gamma^2\mathbf{A}^2\Sigma(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k})}{d(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^2)}. \\
\Rightarrow V_{Lookaround}^* &= \frac{1}{\mathbf{I} - \frac{d-1}{d}(\mathbf{I} - \gamma\mathbf{A})^{2k}} \frac{\gamma^2\mathbf{A}^2\Sigma(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k})}{d(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^2)} \\
&= \frac{\gamma^2\mathbf{A}^2\Sigma[\mathbf{I} - (\mathbf{I} - \gamma\mathbf{I})^{2k}]}{[d\mathbf{I} - (d-1)(\mathbf{I} - \gamma\mathbf{A})^{2k}][\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^2]} \\
&= \frac{\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k}}{d\mathbf{I} - (d-1)(\mathbf{I} - \gamma\mathbf{A})^{2k}}V_{SGD}^*.
\end{aligned}$$

56 According to Equation 6, we can deduce that

$$\begin{aligned}
V_{Lookaround}^* &= \frac{\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k}}{d\mathbf{I} - (d-1)(\mathbf{I} - \gamma\mathbf{A})^{2k}}V_{SGD}^* \\
&= \frac{\alpha^2(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k}) + 2\alpha(1-\alpha)(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^k)}{\alpha^2(d\mathbf{I} - (d-1)(\mathbf{I} - \gamma\mathbf{A})^{2k})}V_{Lookahead}^*.
\end{aligned}$$

57 The proof is now complete. \square

58 B.1 Comparing the dynamics of Lookahead

59 We now proceed with the proof for the range of constraints variable α in Equation 3. When $d \geq 3$,
60 and $\alpha \in [0.5, 1)$, the Lookaround method can obtain a smaller variance than the Lookahead method:

61 *Proof.* Let $B = (\mathbf{I} - \gamma\mathbf{A})^k$, due to $0 < \gamma < 1/L$, $L = \max_i a_i$, so we can have $B \preceq \mathbf{I}$. Substituting
62 the matrix B into the expressions for the variance fixed point relation of Lookaround and Lookahead,
63 we can obtain

$$\begin{aligned}
V_{Lookaround}^* &= \frac{\alpha^2(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^{2k}) + 2\alpha(1-\alpha)(\mathbf{I} - (\mathbf{I} - \gamma\mathbf{A})^k)}{\alpha^2(d\mathbf{I} - (d-1)(\mathbf{I} - \gamma\mathbf{A})^{2k})}V_{Lookahead}^* \\
&= \frac{\mathbf{I} - B^2 + 2\frac{1-\alpha}{\alpha}(\mathbf{I} - B)}{d\mathbf{I} - (d-1)B^2}V_{Lookahead}^* \\
&= \frac{\frac{2-\alpha}{\alpha}\mathbf{I} - B^2 - \frac{2-2\alpha}{\alpha}B}{d\mathbf{I} - (d-1)B^2}V_{Lookahead}^*.
\end{aligned}$$

64 Let the coefficient matrix be denoted as C, when $d \geq 3$, for each diagonal element C_{ii} of C, we can
65 scale the denominator of this expression as follows:

$$C_{ii} \leq \frac{\frac{2-\alpha}{\alpha} - B_{ii}^2 - \frac{2-2\alpha}{\alpha}B_{ii}}{3 - 2B_{ii}^2},$$

66 Then, we can derive the range of α by restricting the right-hand side expression to be less than or
67 equal to 1.

$$\frac{\frac{2-\alpha}{\alpha} - B_{ii}^2 - \frac{2-2\alpha}{\alpha}B_{ii}}{3 - 2B_{ii}^2} \leq 1,$$

68 As $0 \leq B_{ii} \leq 1$, we can multiply both sides of the inequality by the denominator, then rearrange and
69 combine like terms to obtain the following form:

$$B_{ii}^2 - \frac{2-2\alpha}{\alpha}B_{ii} + \frac{2-4\alpha}{\alpha} \leq 0.$$

70 Skipping the detailed steps, we can obtain $\alpha \geq 0.5$ by solving the quadratic equation. Therefore, in
71 the case where $\alpha \in [0.5, 1)$ and $d \geq 3$ ($\alpha < 1$ is subject to Lookahead's settings), the coefficient
72 matrix $C \preceq \mathbf{I}$, so the convergence speed of Lookaround is slower than Lookahead.

73 \square

74 **B.2 Deterministic quadratic convergence**

75 Since our method samples data under multiple data augmentations, it is approximately seen as the
 76 average sampling of historical trajectories in the convergence analysis of quadratic functions. In such
 77 a perspective, we compare the convergence rates of Lookaround and Lookahead.

78 We first show the state transition equation for the classical momentum method in quadratic functions:

$$\mathbf{v}_{t+1} = \beta \mathbf{v}_t - \nabla_{\theta} f(\theta_t) = \beta \mathbf{v}_t - \mathbf{A} \theta_t, \quad (7)$$

$$\theta_{t+1} = \theta_t + \gamma \mathbf{v}_{t+1} = \gamma \beta \mathbf{v}_t + (\mathbf{I} - \gamma \mathbf{A}) \theta_t. \quad (8)$$

79 Here, \mathbf{v} stands for the momentum term. We can generalize this to matrix form:

$$\begin{bmatrix} \mathbf{v}_{t+1} \\ \theta_{t+1} \end{bmatrix} = \begin{bmatrix} \beta & -\mathbf{A} \\ \gamma \beta & \mathbf{I} - \gamma \mathbf{A} \end{bmatrix} \begin{bmatrix} \mathbf{v}_t \\ \theta_t \end{bmatrix}.$$

80 Thus, given the initial θ_0 , we can obtain the convergence rate with respect to θ by the maximum
 81 eigenvalue of the matrix. Referring to Zhang et al. [7] and Lucas et al. [4], we obtain the state
 82 transition matrix regarding our algorithm as follows:

$$\begin{bmatrix} \theta_{t,0} \\ \theta_{t-1,k} \\ \vdots \\ \theta_{t-1,1} \end{bmatrix} = LB^{(k-1)}T \begin{bmatrix} \theta_{t-1,0} \\ \theta_{t-2,k} \\ \vdots \\ \theta_{t-2,1} \end{bmatrix},$$

83 where L, B and T denote the average weight matrix, the single-step transfer matrix and the position
 84 transformation matrix respectively:

$$L = \begin{bmatrix} \frac{1}{k+1}I & \frac{1}{k+1}I & \cdots & \frac{1}{k+1}I & \frac{1}{k+1}I \\ I & 0 & \cdots & \cdots & 0 \\ 0 & I & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots \\ 0 & \cdots & 0 & I & 0 \end{bmatrix},$$

$$B = \begin{bmatrix} (1 + \beta)I - \eta A & -\beta I & 0 & \cdots & 0 \\ I & 0 & \cdots & \cdots & 0 \\ 0 & I & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots \\ 0 & \cdots & 0 & I & 0 \end{bmatrix},$$

$$T = \begin{bmatrix} I - \eta A & \beta I & -\beta I & 0 & \cdots & 0 \\ I & 0 & \cdots & \cdots & 0 & \vdots \\ 0 & I & \ddots & \cdots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots & \vdots \\ \vdots & \cdots & 0 & I & 0 & 0 \\ 0 & \cdots & 0 & 0 & I & 0 \end{bmatrix}.$$

85 After specifying the appropriate parameters and performing matrix multiplication to obtain the state
 86 transition matrix, the convergence rate ρ can be obtained by calculating the largest eigenvalue of the
 87 matrix. Note that since this linear dynamical system corresponds to k updates, we finally have to
 88 compute the k^{th} root of the eigenvalues to recover the correct convergence bounds.

89 C Experimental Detail

90 C.1 Random Initialization

91 C.1.1 CIFAR10 and CIFAR100

92 **Data augmentation details.** For the CIFAR10 dataset, we use [RandomCrop + *] for data aug-
93 mentation, and for the CIFAR100 dataset, we use [RandomCrop + * + RandomRotation] for data
94 augmentation. * can be replaced by three different data augmentation methods of random horizontal
95 flip, random vertical flip, or RandAugment [1].

96 **Training details.** For the CIFAR10 and CIFAR100 datasets, we have applied some common
97 settings. The initial learning rate is set to 0.1, and the batch size is set to 128. Additionally, a warm-up
98 phase of 1 epoch is implemented. Subsequently, different learning rate schedulers are used based on
99 the specific dataset. For the CIFAR100 dataset, we utilize the MultiStepLR scheduler. The learning
100 rate is decayed at the 60, 120, and 160 epochs, with a decay factor of 0.2. For the CIFAR10 dataset,
101 we employ the CosineAnnealingLR learning rate scheduler. In comparison with other optimizers or
102 optimization methods, we use the default settings of the other methods in the corresponding papers.

103 C.1.2 ImageNet.

104 For the ImageNet dataset, we use [RandomResizedCrop + * + RandomRotation] for data augmen-
105 tation. * can be replaced by three different data augmentation methods of random horizontal flip,
106 random vertical flip, or RandAugment [1]. We train the model with the following settings: an initial
107 learning rate of 0.1, a batch size of 256, and a warm-up phase of 1 epoch. We train the model for a
108 total of 90 epochs. We utilize the MultiStepLR scheduler with decay steps at 30 and 60 epochs, and a
109 decay factor of 0.1.

110 C.2 Finetuning

111 In this stage, all images are resized to 224*224 pixels to fit the input size of the pretrained model, all
112 models use the ImageNet-1k pretrained weights from the PyTorch library, and other settings remain
113 the same as in C.1.1. For the training of ViT-B/16, we utilize the Adam optimizer with an initial
114 learning rate of 0.001. β_1 is set to 0.9, β_2 is set to 0.999, and the weight decay is set to 0.00005. To
115 reduce memory consumption, we employed a batch size of 64.

116 C.3 Compared with Ensemble Method

117 We compare Lookaround with Logit Ensemble and Snapshot Ensemble [3]. In the setting of Logit
118 Ensemble, we train multiple models separately using different data augmentation methods, and then
119 average the outputs of these models for prediction. However, this approach requires more inference
120 time. In the setting of Snapshot Ensemble, we use the CosineAnnealingWarmRestarts learning rate
121 scheduler to collect four snapshots during the training process. Then, we average the outputs of these
122 different snapshots for prediction. This approach also requires more inference time.

123 C.4 Ablation Study

124 In ablation experiments, we consider the effectiveness of independent components Data Augmenta-
125 tion (DA) and Weight Averaging (WA) for Lookaround, respectively. In the ablation experiments
126 without data augmentation, in order to improve the competitiveness of the corresponding experiments,
127 we choose to train different models independently using different data augmentation methods, and
128 then select the model that performs best on the test dataset.

129 C.5 Additional Analysis

130 In the robustness experiments with the number of Data Augmentation methods, the six data augmen-
131 tation methods are given as: RandomVerticalFlip, RandomHorizontalFlip, RandAugment, AutoAug-
132 ment, RandomPerspective, RandomEqualize. All the Augmentation methods are from the Pytorch
133 library. When more data augmentation methods are used, the training time will be correspondingly

134 increased in our method. Therefore, in the main experiments in this paper, in order to reduce the time
135 consumption, we only select three data augmentation methods for comparison.

136 D Relationship with Model Soups

137 Model Soups [5] is a framework for finetuning a common pretrained model using different hyperpa-
138 rameters and then averaging the weights of different finetuned models to improve model performance
139 and generalization. We conduct experiments under two data augmentation methods to compare the
140 performance gap between Lookaround and Model Soups.

141 As shown in Figure 2, Model Soups is effective with a few finetuning epochs, achieving better
142 performance than individual models (θ_1 or θ_2). However, as the number of training epochs increases,
143 its robustness significantly decreases, leading to a result that is unacceptable. We can observe that
144 the accuracy after 50 epochs is almost near 0. Different models in Model Soups easily fall into
145 different loss basins and cannot be connected linearly. On the contrary, Lookaround continuously
146 performs weight averaging during training to maintain the locality of different models for linear mode
147 connectivity, resulting in better robustness and generalization performance.

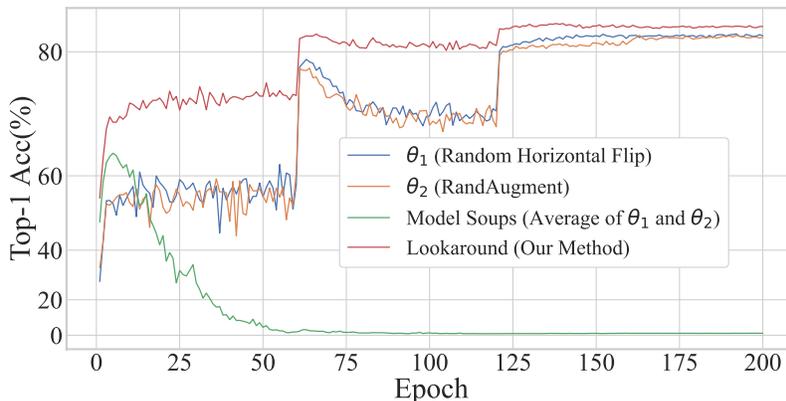


Figure 2: Top-1 accuracy curves of ResNet50 on CIFAR100 under Lookaround and Model Soups. The initial weights of the ResNet50 model are obtained from the pretrained weights on the ImageNet.

148 References

- 149 [1] Ekin D. Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V. Le. Randaugment: Practical automated data
150 augmentation with a reduced search space. *Conference on Neural Information Processing Systems*, 2019.
- 151 [2] Timur Garipov, Pavel Izmailov, D. A. Podoprikin, Dmitry Vetrov, and Andrew Gordon Wilson. Loss
152 surfaces, mode connectivity, and fast ensembling of dnns. *Conference on Neural Information Processing
153 Systems*, 2018.
- 154 [3] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E. Hopcroft, and Kilian Q. Weinberger. Snapshot
155 ensembles: Train 1, get M for free. In *International Conference on Learning Representations*, 2017.
- 156 [4] James Lucas, Shengyang Sun, Richard Zemel, and Roger Grosse. Aggregated momentum: Stability through
157 passive damping. In *International Conference on Learning Representations*, 2018.
- 158 [5] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo Lopes, Ari S.
159 Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model
160 soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference
161 time. In *International Conference on Machine Learning*, 2022.
- 162 [6] Yuhuai Wu, Mengye Ren, Renjie Liao, and Roger Grosse. Understanding short-horizon bias in stochastic
163 meta-optimization. *arXiv preprint arXiv:1803.02021*, 2018.
- 164 [7] Michael R. Zhang, James Lucas, Jimmy Ba, and Geoffrey E. Hinton. Lookahead optimizer: k steps forward,
165 1 step back. *Conference on Neural Information Processing Systems*, 2019.