

467 **A Proof of (3)**

468 Due to fundamental theorem of calculus,

$$L(\theta^{k+1}) - L(\theta^k) = \int_C \nabla L(\theta) \cdot d\theta = \int_0^1 \nabla L(\theta(t)) \cdot v(t) dt, \tag{5}$$

469 where  $C$  is a trajectory whose start and end points are  $\theta^k$  and  $\theta^{k+1}$ . In GD setting, because  $\theta^{k+1} =$   
 470  $\theta^k - \alpha \nabla L(\theta^k)$  for some learning rate  $\alpha > 0$ , we can think of the straight line trajectory joining  $\theta^k$   
 471 and  $\theta^{k+1}$ . In this case, the velocity vector becomes  $v(t) = -\alpha \nabla L(\theta^k)$  and

$$L(\theta^{k+1}) - L(\theta^k) = -\alpha \int_0^1 \nabla L(\theta(t)) \cdot \nabla L(\theta^k) dt \tag{6}$$

472 where  $\theta(0) = \theta^k$ ,  $\theta(1) = \theta^{k+1}$  and  $\theta(t) = (1 - t)\theta(0) + t\theta(1)$ .

473 By Taylor series expansion it becomes

$$\nabla L(\theta(t)) \approx \nabla L(\theta^k) + H(\theta^k)(\theta(t) - \theta^k) = (I - \alpha t H(\theta^k)) \nabla L(\theta^k) \tag{7}$$

474 and combining Eq.(6) and Eq.(7) proves Eq.(3).

475 **B SHOT is robust to hyperparameter settings**

Table 8: Test accuracy % of 4-conv using SHOT with different  $\lambda$  values. The values in parentheses indicate the number of shots. The better accuracy between the baseline and SHOT is bold-faced.

size of $\lambda$	miniImageNet (5)
0 (baseline, MAML)	64.81 $\pm$ 1.63
0.1 (SHOT <sub>r</sub> )	<b>66.86</b> $\pm$ 0.58
10 <sup>-2</sup>	65.97 $\pm$ 0.93
10 <sup>-3</sup>	66.85 $\pm$ 0.47
10 <sup>-4</sup>	66.08 $\pm$ 0.36
10 <sup>-5</sup>	66.81 $\pm$ 0.28

476 Table 8 shows the performance of SHOT with various hyperparameter settings. We conducted  
 477 experiments with different learning rates of  $\lambda$ , and in all cases, SHOT improved the performance of  
 478 the baseline. This suggests that SHOT is robust to hyperparameter settings.

479 **C Another viewpoint of ANIL and BOIL**

480 In this section, we reconcile the opposite opinions of feature reuse versus feature adaptation in  
 481 gradient-based meta-learning (GBML) with our hypothesis that reducing the impact of the Hessian in  
 482 the inner loop can improve performance.

483 ANIL, proposed in [11], argues that feature reuse is key in the inner loop, where the feature remains  
 484 invariant while only the decision boundary is adapted. On the other hand, BOIL, proposed in [12],  
 485 argues that feature adaptation is key in the inner loop, where the feature is adapted while the decision  
 486 boundary remains invariant.

487 To test their arguments, ANIL and BOIL proposed two algorithms. ANIL freezes the encoder and  
 488 only updates the head in the inner loop, while BOIL freezes the head and only updates the encoder in  
 489 the inner loop. The problem is that both algorithm shown good performance thereby both arguments  
 490 look persuasive despite they argue exactly in the opposite ways.

491 Our hypothesis, which suggests that the outer loop implicitly suppresses the Hessian along the  
 492 optimization trajectory, can reconcile the arguments of both ANIL and BOIL. This is because our  
 493 hypothesis implies that the model acts linearly in the inner loop. ANIL and BOIL can be interpreted as  
 494 algorithms that enforce linearity in the inner loop by restricting parameters and reducing the number  
 495 of non-linear components between layers.

496 ANIL freezes the encoder and only updates the head in the inner loop, reducing the number of non-  
 497 linear components in the inner loop. This enforces linearity in the inner loop, as the only non-linearity  
 498 is the loss function. ANIL achieves better performance than MAML in 1-step optimization, as it is  
 499 more powerful at 1-step optimization, which views the model as linear.

500 BOIL freezes the head and only updates the encoder in the inner loop, reducing the number of  
 501 non-linear components in the inner loop. By applying BOIL, the gradient norm is predominant in the  
 502 last layer of the encoder, making it a variant of ANIL that updates only the penultimate layer. This  
 503 layer has much stronger performance, as it can change the feature while maintaining the linearity of  
 504 the model. Table 14 of [12] shows a boosted performance when all but the penultimate layer is not  
 505 frozen. By explicitly enforcing linearity in the inner loop, BOIL achieves improved performance.

## 506 D GBML is a variant of Prototype Vector method

507 In this section, we provide a novel viewpoint of GBML, that GBML (Gradient-Based Meta Learning)  
 508 is a variant of MBML (Metric-Based Meta Learning). This viewpoint relies on the **linearity assumption**.  
 509 *i.e.*, the effect of the Hessian along the optimization trajectory is zero, thereby the model act as  
 510 linear in the inner loop.

511 Suppose there exists a meta-learning model that satisfies the linearity assumption in the inner loop,  
 512 then classifying a new classification task with a task-specific function  $f(\cdot|\theta^*)$  after an inner loop is  
 513 equivalent to creating a prototype vector for each class on a specific feature map and classifying the  
 514 input as the class of the most similar prototype vector.

515 The proof starts by defining the prototype vector at first.

516 **Prototype Vector** We define a prototype vector  $V_c$  for class  $c$  in an  $N$ -way  $K$ -shot classification task  
 517 formally as

$$V_c = \sum_{i=1}^N \sum_{j=1}^K \beta_{ij} \varphi_c(X_{ij}), \quad c \in \{1, \dots, N\}, \quad (8)$$

518 where  $X_{ij}$  is the  $j$ -th input sample for the  $i$ -th class,  $\varphi_c(\cdot) \in \mathcal{H}$  is a class-specific feature map and  
 519  $\beta_{ij}$  indicates the importance of  $X_{ij}$  for constituting the prototype vector  $V_c$ .

520 In other words, there exists a feature map  $\varphi_c$  for each class  $c$ , and the support set is mapped to the  
 521 corresponding feature map and then weighted-averaged to constitute the prototype vector of the  
 522 corresponding class. At inference time, the classification of a given query  $X$  is done by taking the  
 523 class of the most similar prototype vector as follows:

$$\hat{c} = \arg \max_c \langle V_c, \varphi(X) \rangle. \quad (9)$$

524 Here,  $\varphi : \mathcal{X} \rightarrow H$  is a non-class-specific mapping. We can also rewrite the prototype vector using  $\varphi$   
 525 and by defining a projection  $P_c : \mathcal{X} \rightarrow \mathcal{X}$  as

$$P_c(X) = \begin{cases} X & \text{if } y(X) = c, \\ \nu \in \mathcal{N}(\varphi), & \text{if } y(X) \neq c \end{cases} \quad (10)$$

526 where  $y(X)$  is the ground truth class of  $X$  and  $\mathcal{N}$  is the null space of  $\varphi$  *i.e.*,  $\varphi(\nu) = 0$ .

527 Then by defining  $\varphi_c \triangleq \varphi \circ P_c$  and  $\beta_{ij} \triangleq \frac{1}{K}$ , it becomes

$$V_c = \frac{1}{K} \sum_{j=1}^K \varphi(X_{cj}). \quad (11)$$

528 **SGD in the inner loop** If GBML satisfies the hypothesis of linearity in the inner loop,  $f$  is locally  
 529 linear in  $\theta$  in an inner loop. More specifically, there exists an equivalent feature map  $\varphi_c : \mathcal{X} \rightarrow \mathcal{H}$   
 530 which satisfies  $f_c(\cdot|\theta_c) = \langle \theta_c, \varphi_c(\cdot) \rangle$  for every  $x \in \mathcal{X}$  where  $f(\cdot|\theta) = [f_1(\cdot|\theta_1), \dots, f_N(\cdot|\theta_N)]^T$ .

531 With the loss function  $L(x, y|\theta) = D(s(f(x|\theta)), y)$  for some distance measure  $D$  such as cross  
 532 entropy, we can formulate the inner loop of  $N$ -way  $K$ -shot meta learning by SGD as

$$\theta_c^{k+1} = \theta_c^k - \alpha \sum_{i=1}^N \sum_{j=1}^K \frac{\partial L(X_{ij}, y(X_{ij})|\theta)}{\partial \theta_c} = \theta_c^k - \alpha \sum_{i=1}^N \sum_{j=1}^K \frac{\partial D}{\partial f_c} \varphi_c(X_{ij}), \quad (12)$$

533 since all samples in the support set are inputted in a batch of an inner loop.

534 Because the model is linear in the inner loop, the batch gradient does not change. Let  $\beta_{ij} =$   
 535  $-\frac{\partial D}{\partial f_c}|_{\theta_c^0, X_{ij}}$ . Then after  $t$  steps, by (8), the model becomes

$$\theta_c^t = \theta_c^0 + \alpha t \sum_{i=1}^N \sum_{j=1}^K \beta_{ij} \varphi_c(X_{ij}) = \theta_c^0 + \alpha t V_c. \quad (13)$$

536 At the initialization step of an inner loop, there is no information about the class, even the configuration  
 537 order of the class, because the task is randomly sampled. If so, the problem is solved in the inner loop.  
 538 For example, if a class *Dog* is allocated to a specific index such as *Class 3*. There is no guarantee that  
 539 it will have the identical index the next time the class *Dog* comes in. Thus, at a meta-initialization  
 540 point  $\theta^0$ , the scores for different classes would not be much different, *i.e.*,  $f_i(x|\theta_i^0) \simeq f_j(x|\theta_j^0)$  for  
 541  $i, j \in [1, \dots, N]$ .

542 Considering the goal of classification is achieved through relative values between  $f_i(X)$ 's, the value  
 543 at the initialization point does not need to be considered significantly. Therefore

$$\arg \max_c f_c(X) = \arg \max_c \langle \theta_c^t, \varphi(X) \rangle = \arg \max_c \langle \theta_c^0 + \alpha t V_c, \varphi(X) \rangle \sim \arg \max_c \langle V_c, \varphi(X) \rangle \quad (14)$$

So inner loop in GBML can be interpreted as making proptotype vector with given support set.  $\square$

Table 9: Test accuracy % of 4-conv network on benchmark data sets. The values in parentheses indicate the number of shots. The best accuracy among different methods is bold-faced. To differentiate the notation, we have denoted SHOT<sub>3</sub> as a model that uses 3 optimization steps in  $f_r$  and 1 optimization step in  $f_t$ , and SHOT<sub>6</sub> as a model that uses 6 optimization steps in  $f_r$  and 3 optimization steps in  $f_t$ .

meta-train	miniImageNet			Cars			
meta-test	miniImageNet	tieredImageNet	Cars	Cars	miniImageNet	CUB	
MAML (1)	47.88 ± 0.55	<b>51.84</b> ± 0.24	34.41 ± 0.47	47.78 ± 0.99	28.67 ± 1.17	30.95 ± 1.41	
MAML + SHOT <sub>3</sub> (1)	47.97 ± 0.71	51.68 ± 0.68	34.03 ± 1.11	<b>50.44</b> ± 0.62	<b>30.19</b> ± 0.50	<b>31.21</b> ± 0.64	
MAML + SHOT <sub>6</sub> (1)	<b>48.15</b> ± 0.31	51.67 ± 0.73	<b>34.79</b> ± 0.70	49.89 ± 0.17	28.39 ± 0.37	30.83 ± 0.26	
MAML (5)	64.81 ± 1.63	67.96 ± 1.22	46.57 ± 0.53	62.24 ± 2.01	37.23 ± 1.95	41.84 ± 1.25	
MAML + SHOT <sub>3</sub> (5)	66.86 ± 0.58	69.48 ± 0.17	<b>48.42</b> ± 0.72	<b>69.08</b> ± 0.31	<b>40.79</b> ± 0.93	<b>43.46</b> ± 0.89	
MAML + SHOT <sub>6</sub> (5)	66.27 ± 0.23	<b>69.60</b> ± 0.31	45.83 ± 1.82	66.37 ± 1.93	39.35 ± 0.74	42.63 ± 0.24	

544

## 545 E SHOT with more optimization step in the inner loop

546 In the main paper, we used only one step for the target model to improve computation efficiency.  
 547 However, it's also important to test if SHOT works with more optimization steps in the inner loop. As  
 548 a reference model, we set the number of optimization steps to 6, which is different from the main  
 549 paper where we only used 3 steps (same as the baseline). As shown in Table 9, SHOT still performs  
 550 better than the baseline even with more optimization steps in the inner loop.