

## Problem Setup

#### Algorithm 1: SGD





#### Key steps in SGD & LookAhead (LA):

- 1) inner-loop optimization: K steps forward in SGD & LA
- 2) outer-loop optimization: 1 step back in LA, while no step back in SGD

Optimizer	CIFAR-10	CIFAR-100
SGD	$95.23 \pm .19$	$78.24\pm.18$
Polyak	$95.26 \pm .04$	$77.99 \pm .42$
Adam	$94.84 \pm .16$	$76.88 \pm .39$
LOOKAHEAD	$95.27\pm.06$	$78.34 \pm .05$

Optimizer	TRAIN	VAL.	Test
SGD	43.62	66.0	63.90
LA(SGD)	35.02	65.10	63.04
Adam	33.54	61.64	59.33
LA(ADAM)	31.92	60.28	57.72
Polyak	-	61.18	58.79

#### **ResNet 18**

Important observations: LookAhead (LA) enjoys better test performance than SGD

#### **Problem**:

1) Why LA enjoys better test performance than SGD?

2) How to further improve LA?

**Tools for Test Performance Analysis** 

**Optimal solution to empirical risk** on dataset S:

$$\boldsymbol{\theta}_{\mathcal{S}}^* \in \operatorname{argmin}_{\boldsymbol{\theta}} F_{\mathcal{S}}(\boldsymbol{\theta}) \triangleq \frac{1}{n} \sum_{i=1}^n \ell(f(\boldsymbol{x}_i; \boldsymbol{\theta}),$$

Approximate solution to empirical risk when using algorithm  $\mathcal{A}$  on dataset  $\mathcal{S}$ :

$$\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \approx \underset{\boldsymbol{\theta}}{\operatorname{argmin}} F_{\mathcal{S}}(\boldsymbol{\theta}) \triangleq \frac{1}{n} \sum_{i=1}^{n} \ell(f(\boldsymbol{x}_{i};\boldsymbol{\theta}))$$

**Excess risk error to measure test performance:** 

$arepsilon_{exc} = \mathbb{E}_{\mathcal{A},\mathcal{S}}[F(oldsymbol{ heta}_{\mathcal{A},\mathcal{S}})]$	$ -\mathbb{E}_{\mathcal{A},\mathcal{S}}[F_{\mathcal{S}}(oldsymbol{ heta}_{\mathcal{S}})] =$	$\mathbb{E}_{\mathcal{A},\mathcal{S}}[F(\theta_{\mathcal{A},\mathcal{S}}) - F_{\mathcal{S}}(\theta_{\mathcal{A},\mathcal{S}})]$
test error	best training error	generalization error
$\Gamma(0) \land \Pi$	$\begin{bmatrix} 0 (f(x, 0), x) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$	

where  $F(\theta) \triangleq \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim \mathcal{D}}[\ell(f(\boldsymbol{x};\theta),\boldsymbol{y})]$  is the population risk.

# **Necessary definitions:**

 $\lambda$ -strongly convex:  $\forall \theta_1, \theta_2, f(\theta_1) \ge f(\theta_2) + \langle \nabla f(\theta_2), \theta_1 - \theta_2 \rangle + \frac{\lambda}{2} \| \theta_1 - \theta_2 \|^2$ **CONVEX**:  $\forall \theta_1, \theta_2, f(\theta_1) \ge f(\theta_2) + \langle \nabla f(\theta_2), \theta_1 - \theta_2 \rangle$ *G*-Lipschitz continuous:  $||f(\theta_1) - f(\theta_2)||_2 \le G||\theta_1 - \theta_2||_2$ *L*-smooth:  $\|\nabla f(\theta_1) - \nabla f(\theta_2)\|_2 \leq L \|\theta_1 - \theta_2\|_2 (\forall \theta_1, \theta_2)$ **Polyak-Łojasiewicz (PŁ) Condition**:  $2\mu(f(\theta) - f(\theta^*)) \le \|\nabla f(\theta)\|^2 (\forall \theta)$  where  $\theta^* \in \operatorname{argmin}_{\theta} f(\theta)$ . Weakly Quasi-Convexity:  $\langle \nabla f(\theta), \theta - \theta^* \rangle \ge \rho(f(\theta) - f(\theta^*))$ 

# Towards Understanding Why Lookahead Generalizes Better Than SGD and Beyond

Pan Zhou\*, Hanshu Yan\*, Xiao-Tong Yuan<sup>†</sup>, Jiashi Feng\*, Shuicheng Yan\* <sup>†</sup> Nanjing University of Information Science & Technology, China Sea Al Lab, Singapore {zhoupan, yanhanshu, fengjs, yansc}@sea.com xtyuan@nuist.edu.cn

**Input** :Objective  $F_{\mathcal{S}}(\boldsymbol{\theta})$ , dataset  $\mathcal{S}$ , inner-loop optimizer  $\mathcal{A}$ , inner-loop step number k and learning rate  $\{\{\eta_{\tau}^{(t)}\}\}$ , outer-loop learning rate  $\alpha \in (0, 1)$ .

> Inner-loop optimization  $\boldsymbol{v}_{\tau}^{(t)} = \mathcal{A}(F_{\mathcal{S}}(\boldsymbol{\theta}), \boldsymbol{v}_{\tau-1}^{(t)}, \eta_{\tau-1}^{(t)}, \mathcal{S}) = \boldsymbol{v}_{\tau-1}^{(t)} - \eta_{\tau-1}^{(t)} \boldsymbol{g}_{\tau-1}^{(t)}$

> > outer-loop optimization

LSTM

), **y**<sub>i</sub>)

 $+\mathbb{E}_{\mathcal{A},\mathcal{S}}[F_{\mathcal{S}}(\theta_{\mathcal{A},\mathcal{S}})-F_{\mathcal{S}}(\theta_{\mathcal{S}}^{*})]$ optimization error

#### Main Results

### **Excess risk error to measure test performance:**

test error

best training error

 $\eta = \frac{1}{\sqrt{kT}}$ , on convex problem we have

learning rate  $\eta = \frac{1}{\sqrt{kT}}$ , on convex problem we have

$$\frac{1}{\Gamma^{\alpha}} > \mathcal{O}\left(\frac{1}{n\lambda}\frac{\ln(IK)}{k^{\alpha}}\right)$$

the optimum  $\alpha$  is not 1.

optimization error  $\leq O$ generalization error  $\leq O$  (

where  $\gamma = (1 - \frac{1}{n})\frac{\alpha L}{\mu}$  and  $\xi = \ell_{\max}^{\frac{1}{1+\gamma}}$ 

**Remark:** with properly  $\alpha$ , **lookahead may have smaller test error than SGD** 

An Improved LookAhead: Stagewise Locally-regularized Lookahead

**Input** :Loss  $F_{\mathcal{S}}(\boldsymbol{\theta})$ , constant for q = 1, 2, ..., Q do  $F_q(\boldsymbol{\theta}) = F_{\mathcal{S}}(\boldsymbol{\theta}) + \frac{\beta_q}{2} \| \boldsymbol{\theta} \|$  $\boldsymbol{\theta}_q = \text{Look-ahead}(F_q(\boldsymbol{\theta}))$ end **Output** :  $\theta_{\mathcal{A},\mathcal{S}} = \theta_Q$ .



$$\frac{1}{(Tk+1)^{2\alpha}} + \frac{2\alpha LG \left(\alpha + 2(1-\alpha)(K-1)\right)}{\mu^2 (Tk+1)^{2\alpha-1}}\right),$$

$$\frac{\xi}{n-1} \alpha^{\frac{1}{1+\gamma}} (Tk)^{\frac{\gamma}{\gamma+1}}\right).$$

$$\left[\frac{2G^2}{\mu}\right]^{\frac{1}{1+\gamma}} \text{ in which } \ell_{\max} = \max_{\theta, (\boldsymbol{x}, \boldsymbol{y})} \ell(f(\boldsymbol{x}; \theta), \boldsymbol{y}).$$

# Algorithm 3: Stagewise Locally-Regularized LookAhead (SLRLA)

$$\tan\{\beta_q\}_{q=1}^Q$$

$$(\boldsymbol{\theta}_{q-1} \|^2;$$
  
 $(\boldsymbol{\theta}_{q}), \eta_q, T_q, \alpha_q, k_q, \boldsymbol{\theta}_{q-1}, \mathcal{A}, \mathcal{S}).$   
Vanilla LookAhead

# An Improved LookAhead: Stagewise Locally-regularized Lookahead

Strategy: divide optimization into several stages and use lookahead to solve locally-regularized loss

# Advantages:

the stochastic gradient complexity (stochastic gradient evaluation number, a.k.a. IFO) is

stochastic gra

 $\lambda$ -stronglynonconvex pr

C	Generaliza					
е	rror					
tł	ne generali					
	generaliza					
	$\lambda$ -strongly-co					
	nonconvex prob					

# **Remark**:

Experiments



# **Remark:** SLRLA has better test performance than stagewise LA (SLA)



$$F_q(\theta) = F_S(\theta) + rac{eta_q}{2} \| heta - heta_{q-1} \|_2^2$$

 $\blacktriangleright$  Local regularization improves loss convexity, e.g. ill-conditioned loss —; well-conditioned one

Local regularization helps avoid overfitting

**Optimization error**. Under proper assumptions, to obtain optimization error optimization error  $\leq \epsilon$ ,

adient complexity	LookAhe	SLRLA		
defent complexity	$\alpha \in (0, \frac{1}{2})$	$\alpha = \frac{1}{2}$	$\alpha \in \left(\frac{1}{2}, 1\right]$	$\alpha \in (0,1]$
convex problems	$ \left  \mathcal{O}\left( \left(\frac{1}{\epsilon}\right)^{\frac{1}{2\alpha}} + \left(\frac{1}{(1-2\alpha)\lambda^{2}\epsilon}\right)^{\frac{1}{2\alpha}} \right) \right  $	$\mathcal{O}\left(\frac{\log \frac{1}{\epsilon}}{\lambda^2 \epsilon}\right)$	$\mathcal{O}\left(\frac{1}{(2\alpha-1)\lambda^2\epsilon}\right)$	$\mathcal{O}(\frac{1}{\lambda \alpha \epsilon})$
roblems with $\mu$ -PL	$\mathcal{O}((\frac{1}{\mu^2\epsilon}))$	$\left(\right)^{1/\alpha}$		$\mathcal{O}\left(\frac{1}{\mu\alpha\epsilon}\right)$

**Remark**: By observing factors  $\alpha$ ,  $\lambda$  and  $\mu$ , SLRLA has smaller computational complexity than LA, meaning

### SLRLA has smaller optimization error than LA under a given computational budget

#### ation error. Under proper assumptions, to obtain optimization

optimization error  $\leq \epsilon$ ,

lization error is

ation error	LookAhead (LA) $\alpha \in (0, 1]$	$  \qquad SLRLA \ \alpha \in (0,1]$
onvex problems	$\mathcal{O}\left(\frac{1}{n\lambda}\right)$	$\mathcal{O}\left(\frac{1}{n(\beta/\alpha+\lambda)}\right)$
blems with $\mu$ -PL	$\mathcal{O}\left(\frac{1}{n}(Tk)\frac{\gamma}{\gamma+1}\right) \ (\gamma = \left(1 - \frac{1}{n}\right)\frac{\alpha L}{\mu})$	$\mathcal{O}\left(\frac{1}{n} / \left(\frac{c}{\alpha} + \mu\right)\right) (c \ge 0)$

# **SLRLA** has smaller generalization error than LA

Table 3: Classification accuracy (%).  $^{\circ}$ , \*, †, ‡ are respectively reported in [1], [15], [49], [50].

	ResNet18	CIFAR10 VGG16	WRN-16-10	ResNet18	CIFAR100 VGG16	WRN-16-10	ImageNet ResNet18
	94.84 <sup>\$</sup>	91.08	93.54	76.88 <sup>\$</sup>	64.07	74.81	66.54*
	92.56	91.35	91.68	71.43	64.74	71.64	68.13†
	93.85	90.84	94.16	74.30	63.99	75.92	67.62*
	94.95	90.75	95.95	77.30	63.40	79.63	67.93†
	95.20 <sup>‡</sup>	92.25	95.71	77.02 <sup>‡</sup>	68.63	77.93	70.08‡
]	95.23±0.19 <sup>◊</sup>	92.13±0.02	95.51±0.02	78.24±0.18 <sup>◊</sup>	69.97±0.02	$78.95 \pm 0.03$	70.23 <sup>†</sup>
	95.27±0.06 <sup>◊</sup>	92.38±0.02	95.73±0.02	78.34±0.05 <sup>◊</sup>	70.20±0.04	$79.54 \pm 0.02$	70.30±0.09
	<b>95.47</b> ±0.20	<b>92.63</b> ±0.03	<b>96.08</b> ±0.07	<b>78.58</b> ±0.15	<b>70.63</b> ±0.02	<b>79.85</b> \pm 0.05	<b>70.47</b> ±0.12