# Tracking the gradients using the Hessian: A new look at variance reducing stochastic methods 

## Robert M. Gower

Joint work with Nicolas Le Roux and Francis Bach


Research at Google

AISTATS 2018
Playa Blanca, Lanzarote, Canary Islands, April 9-11, 2018

## Solve Empirical Risk Minimization

$$
\min _{\theta \in \mathbf{R}^{d}} f(\theta) \stackrel{\text { def }}{=} \frac{1}{n} \sum_{i=1}^{n} f_{i}(\theta)
$$

where $n$ is the num of data points and $d$ the num of features.

Datum functions

$$
f_{i}(\theta) \text { is twice differentiable }
$$

Ridge Regression

$$
f_{i}(\theta)=\left(y^{i}-\left\langle\theta, x^{i}\right\rangle\right)^{2}+\lambda\|\theta\|_{2}^{2}
$$

Logistic regression

$$
f_{i}(\theta)=\ln \left(1+e^{-y^{i}\left\langle\theta, x^{i}\right\rangle}\right)+\lambda\|\theta\|_{2}^{2}
$$

Some neural nets

$$
f_{i}(\theta)=\ldots
$$

## Using a first order gradient method

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

$$
\mathbb{E}\left[g_{t}\right]=\nabla f\left(\theta_{t}\right)
$$

## Using a first order gradient method

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Unbiased

$$
\mathbb{E}\left[g_{t}\right]=\nabla f\left(\theta_{t}\right)
$$

EXE: Stochastic Gradient descent (SGD)

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right), \quad \text { where } i \sim \mathcal{U}\{1, \ldots, n\}
$$

## Stochastic Gradient Descent $\gamma=0.2$



## Using a first order gradient method

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Unbiased

$$
\mathbb{E}\left[g_{t}\right]=\nabla f\left(\theta_{t}\right)
$$

EXE: Stochastic Gradient descent (SGD)

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right), \quad \text { where } i \sim \mathcal{U}\{1, \ldots, n\}
$$

## Using a first order gradient method

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

$$
\mathbb{E}\left[g_{t}\right]=\nabla f\left(\theta_{t}\right)
$$

EXE: Stochastic Gradient descent (SGD)

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right), \quad \text { where } i \sim \mathcal{U}\{1, \ldots, n\}
$$

EXE: SGD with covariates

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}, \text { where } i \sim \mathcal{U}\{1, \ldots, n\}
$$

$$
z_{i} \in \mathbb{R}^{d}, \text { for } i=1, \ldots, n
$$

## Choosing the covariates

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}
$$

1) Correlated to the stochastic gradients

$$
\text { If } \nabla f_{i}\left(\theta_{t}\right) \approx z_{i} \text { then } \mathbb{V A} \mathbb{R}\left(g_{t}\right) \leq \mathbb{V} \mathbb{A} \mathbb{R}\left(\nabla f_{i}\left(\theta_{t}\right)\right)
$$

## Choosing the covariates

SGD with covariates:

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}
$$

1) Correlated to the stochastic gradients

$$
\text { If } \nabla f_{i}\left(\theta_{t}\right) \approx z_{i} \text { then } \mathbb{V A} \mathbb{R}\left(g_{t}\right) \leq \mathbb{V} \mathbb{A} \mathbb{R}\left(\nabla f_{i}\left(\theta_{t}\right)\right)
$$

2) Cheap to compute

$$
\operatorname{cost}\left(g_{t}\right) \leq \operatorname{cost}\left(\frac{1}{n} \sum_{j=1}^{n} \nabla f_{j}\left(\theta_{t}\right)\right)
$$

## Choosing the covariates

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}
$$

1) Correlated to the stochastic gradients

$$
\text { If } \nabla f_{i}\left(\theta_{t}\right) \approx z_{i} \text { then } \mathbb{V A} \mathbb{R}\left(g_{t}\right) \leq \mathbb{V} \mathbb{A} \mathbb{R}\left(\nabla f_{i}\left(\theta_{t}\right)\right)
$$

2) Cheap to compute

$$
\operatorname{cost}\left(g_{t}\right) \leq \operatorname{cost}\left(\frac{1}{n} \sum_{j=1}^{n} \nabla f_{j}\left(\theta_{t}\right)\right)
$$

EXE: Too costly

$$
\begin{aligned}
z_{i} & =\nabla f_{i}\left(\theta_{t}\right) \\
g_{t} & =\nabla f\left(\theta_{t}\right)
\end{aligned}
$$

## Choosing the covariates

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}
$$

1) Correlated to the stochastic gradients

$$
\text { If } \nabla f_{i}\left(\theta_{t}\right) \approx z_{i} \text { then } \mathbb{V A} \mathbb{R}\left(g_{t}\right) \leq \mathbb{V} \mathbb{A} \mathbb{R}\left(\nabla f_{i}\left(\theta_{t}\right)\right)
$$

2) Cheap to compute

$$
\operatorname{cost}\left(g_{t}\right) \leq \operatorname{cost}\left(\frac{1}{n} \sum_{j=1}^{n} \nabla f_{j}\left(\theta_{t}\right)\right)
$$

EXE: Too costly

$$
\begin{aligned}
& z_{i}=\nabla f_{i}\left(\theta_{t}\right) \\
& g_{t}=\nabla f\left(\theta_{t}\right)
\end{aligned}
$$

EXE: High variance

$$
\begin{aligned}
z_{i} & =0 \\
g_{t} & =\nabla f_{i}\left(\theta_{t}\right)
\end{aligned}
$$

## Choosing the covariates

SGD with covariates:

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1}^{n} z_{j}
$$

1) Correlated to the stochastic gradients

$$
\text { If } \nabla f_{i}\left(\theta_{t}\right) \approx z_{i} \text { then } \mathbb{V A} \mathbb{R}\left(g_{t}\right) \leq \mathbb{V} \mathbb{A} \mathbb{R}\left(\nabla f_{i}\left(\theta_{t}\right)\right)
$$

2) Cheap to compute

$$
\operatorname{cost}\left(g_{t}\right) \leq \operatorname{cost}\left(\frac{1}{n} \sum_{j=1}^{n} \nabla f_{j}\left(\theta_{t}\right)\right)
$$

EXE: Too costly

$$
\begin{aligned}
& z_{i}=\nabla f_{i}\left(\theta_{t}\right) \\
& g_{t}=\nabla f\left(\theta_{t}\right)
\end{aligned}
$$

EXE: High variance

$$
\begin{aligned}
z_{i} & =0 \\
g_{t} & =\nabla f_{i}\left(\theta_{t}\right)
\end{aligned}
$$

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

$$
\tilde{\theta} \in \mathbb{R}^{d}
$$

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

Sample
$\tilde{\theta} \in \mathbb{R}^{d}$
$\nabla f_{i}\left(\theta_{t}\right), \quad i \in\{1, \ldots, n\}$ uniformly

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point
$\tilde{\theta} \in \mathbb{R}^{d}$
Sample
$\nabla f_{i}\left(\theta_{t}\right), \quad i \in\{1, \ldots, n\}$ uniformly
0th order
Taylor

$$
\left\|\tilde{\theta}-\theta_{t}\right\| \text { is small } \Rightarrow \nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})
$$

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point
$\tilde{\theta} \in \mathbb{R}^{d}$
Sample
$\nabla f_{i}\left(\theta_{t}\right), \quad i \in\{1, \ldots, n\}$ uniformly
0th order
Taylor

$$
\left\|\tilde{\theta}-\theta_{t}\right\| \text { is small } \Rightarrow \nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})=: z_{i}
$$

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point
$\tilde{\theta} \in \mathbb{R}^{d}$
Sample
$\nabla f_{i}\left(\theta_{t}\right), \quad i \in\{1, \ldots, n\}$ uniformly
0th order
Taylor
SVRG

$$
\begin{aligned}
& \left\|\tilde{\theta}-\theta_{t}\right\| \text { is small } \Rightarrow \nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})=: z_{i} \\
& g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta})
\end{aligned}
$$

## SVRG: Stochastic Variance Reduced Gradients

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point
$\tilde{\theta} \in \mathbb{R}^{d}$
Sample
$\nabla f_{i}\left(\theta_{t}\right), \quad i \in\{1, \ldots, n\}$ uniformly
0th order
Taylor
SVRG

$$
\begin{aligned}
& \left\|\tilde{\theta}-\theta_{t}\right\| \text { is small } \Rightarrow \nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})=: z_{i} \\
& g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& g_{t}=\nabla f_{i}\left(\theta_{t}\right)-z_{i} \quad+\frac{1}{n} \sum_{j=1} z_{j}
\end{aligned}
$$

## SVRG: Stochastic Variance Reduced <br> Gradients

Set $\theta_{0}=0$, choose $\gamma>0, m \in \mathbb{N}$
$\tilde{\theta}_{0}=\theta_{0}$
for $k=0,1,2, \ldots, T-1$
calculate $\nabla f\left(\tilde{\theta}_{k}\right)$
$\theta_{0}=\tilde{\theta}_{k}$
for $t=0,1,2, \ldots, m-1$
sample $i \in\{1, \ldots, n\}$
$g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}\left(\tilde{\theta}_{k}\right)+\nabla f\left(\tilde{\theta}_{k}\right)$
$\theta_{t+1}=\theta_{t}-\gamma g_{t}$
$\tilde{\theta}_{k+1}=\theta_{m}$
Output $\tilde{\theta}_{T}$

## SVRG: Stochastic Variance Reduced <br> Gradients

$$
\begin{aligned}
& \hline \text { Set } \theta_{0}=0, \text { choose } \gamma>0, m \in \mathbb{N} \\
& \tilde{\theta}_{0}=\theta_{0}
\end{aligned}
$$

$$
\text { for } k=0,1,2, \ldots, T-1
$$

$$
\text { calculate } \nabla f\left(\tilde{\theta}_{k}\right)
$$

$$
\theta_{0}=\tilde{\theta}_{k}
$$

$$
\text { for } t=0,1,2, \ldots, m-1
$$

$$
\text { sample } i \in\{1, \ldots, n\}
$$

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}\left(\tilde{\theta}_{k}\right)+\nabla f\left(\tilde{\theta}_{k}\right)
$$

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

$$
\tilde{\theta}_{k+1}=\theta_{m}
$$

Output $\tilde{\theta}_{T}$

## SVRG: Stochastic Variance Reduced <br> Gradients

> | Set $\theta_{0}=0$, choose $\gamma>0, m \in \mathbb{N}$ |
| :--- |
| $\tilde{\theta}_{0}=\theta_{0}$ |

for $k=0,1,2, \ldots, T-1$

$$
\text { calculate } \nabla f\left(\tilde{\theta}_{k}\right)
$$

$$
\theta_{0}=\tilde{\theta}_{k}
$$

$$
\text { for } t=0,1,2, \ldots, m-1
$$

$$
\text { sample } i \in\{1, \ldots, n\}
$$

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}\left(\tilde{\theta}_{k}\right)+\nabla f\left(\tilde{\theta}_{k}\right)
$$

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

$$
\tilde{\theta}_{k+1}=\theta_{m}
$$

Why not
$1^{\text {st }}$ Taylor?

Output $\tilde{\theta}_{T}$

## SVRG2: Second order tracking

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

1st order
Taylor exp.
$\tilde{\theta} \in \mathbb{R}^{d}$

$$
\nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})+H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right)
$$

## SVRG2: Second order tracking

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

$$
\begin{gathered}
\tilde{\theta} \in \mathbb{R}^{d} \quad H_{i}(\tilde{\theta}):=\nabla^{2} f_{i}(\tilde{\theta}) \\
\nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})+H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \quad=: z_{i}
\end{gathered}
$$

## SVRG2: Second order tracking

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

$$
\begin{gathered}
\tilde{\theta} \in \mathbb{R}^{d} \quad H_{i}(\tilde{\theta}):=\nabla^{2} f_{i}(\tilde{\theta}) \\
\nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})+H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right)=: z_{i} \\
\frac{1}{n} \sum_{j=1} z_{j}=\nabla f(\tilde{\theta})+\frac{1}{n} \sum_{i=1} H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right)
\end{gathered}
$$

## SVRG2: Second order tracking

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

$$
\tilde{\theta} \in \mathbb{R}^{d}
$$

$$
H_{i}(\tilde{\theta}):=\nabla^{2} f_{i}(\tilde{\theta})
$$

$$
\nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})+H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \quad=: z_{i}
$$

$$
\frac{1}{n} \sum_{j=1} z_{j}=\nabla f(\tilde{\theta})+\frac{1}{n} \sum_{i=1} H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right)
$$

SVRG2

$$
\begin{aligned}
g_{t}= & \nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1} z_{j} \\
=\nabla & f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& \quad+\left(\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta})-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right)
\end{aligned}
$$

## SVRG2: Second order tracking

$$
\theta_{t+1}=\theta_{t}-\gamma g_{t}
$$

Reference point

$$
\tilde{\theta} \in \mathbb{R}^{d}
$$

$$
H_{i}(\tilde{\theta}):=\nabla^{2} f_{i}(\tilde{\theta})
$$

$$
\nabla f_{i}\left(\theta_{t}\right) \approx \nabla f_{i}(\tilde{\theta})+H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \quad=: z_{i}
$$

$$
\frac{1}{n} \sum_{j=1} z_{j}=\nabla f(\tilde{\theta})+\frac{1}{n} \sum_{i=1} H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right)
$$

SVRG2

$$
\begin{aligned}
g_{t}= & \nabla f_{i}\left(\theta_{t}\right)-z_{i}+\frac{1}{n} \sum_{j=1} z_{j} \\
=\nabla & \nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& \quad+\left(\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta})-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right)
\end{aligned}
$$

## SVRG2: Stochastic Variance Reduced Gradients with tracking

Set $\theta_{0}=0$, choose $\gamma>0, m \in \mathbb{N}$ $\tilde{\theta}=\theta_{0}$
for $k=0,1,2, \ldots, T-1$
calculate $\nabla f(\tilde{\theta}), H=\nabla^{2} f(\tilde{\theta})$

$$
\theta_{0}=\tilde{\theta}
$$

$$
\text { for } t=0,1,2, \ldots, m-1
$$

$$
\text { sample } i \in\{1, \ldots, n\}
$$

$$
g_{t}=\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta})
$$

$$
\begin{aligned}
& +\left(H-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right) \\
& =\theta_{t}-\gamma g_{t}
\end{aligned}
$$

$$
\tilde{\theta}=\theta_{m}
$$

Output $\tilde{\theta}$

## SVRG2: first experiment



## SVRG2: first experiment



## SVRG2: Stochastic Variance Reduced Gradients with tracking

Set $\theta_{0}=0$, choose $\gamma>0, m \in \mathbb{N}$ $\tilde{\theta}=\theta_{0}$
for $k=0,1,2, \ldots, T-1$ calculate $\nabla f(\tilde{\theta}), H=\nabla^{2} f(\tilde{\theta})$

$$
\theta_{0}=\tilde{\theta}
$$

$$
\text { for } t=0,1,2, \ldots, m-1
$$

$$
\text { sample } i \in\{1, \ldots, n\}
$$

$$
\begin{aligned}
g_{t} & =\nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& +\left(H-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right) \\
\theta_{t+1} & =\theta_{t}-\gamma g_{t}
\end{aligned}
$$

$$
\tilde{\theta}=\theta_{m}
$$

Output $\tilde{\theta}$

## Cost of SVRG2

$$
\begin{aligned}
g_{t}= & \nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& +\left(\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta})-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right)
\end{aligned}
$$

- Full Hessian $H=\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta}) \operatorname{costs} O\left(n d \times \operatorname{eval}\left(f_{i}\right)\right)$
- Hessian vector product $H\left(\theta_{t}-\tilde{\theta}\right) \operatorname{costs} O\left(d^{2}\right)$
- Directional derivative $H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \operatorname{costs} O\left(\operatorname{eval}\left(f_{i}\right)\right)$


## Cost of SVRG2

$$
\begin{aligned}
g_{t}= & \nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& +\left(\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta})-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right)
\end{aligned}
$$

- Full Hessian $H=\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta}) \operatorname{costs} O\left(n d \times \operatorname{eval}\left(f_{i}\right)\right)$
- Hessian vector product $H\left(\theta_{t}-\tilde{\theta}\right)$ costs $O\left(d^{2}\right)$
- Directional derivative $H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \operatorname{costs} O\left(\operatorname{eval}\left(f_{i}\right)\right)$


## Cost of SVRG2

$$
\begin{aligned}
g_{t}= & \nabla f_{i}\left(\theta_{t}\right)-\nabla f_{i}(\tilde{\theta})+\nabla f(\tilde{\theta}) \\
& +\left(\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta})-H_{i}(\tilde{\theta})\right)\left(\theta_{t}-\tilde{\theta}\right)
\end{aligned}
$$

- Full Hessian $H=\frac{1}{n} \sum_{j=1}^{n} H_{j}(\tilde{\theta}) \operatorname{costs} O\left(n d \times \operatorname{eval}\left(f_{i}\right)\right)$
- Hessian vector product $H\left(\theta_{t}-\tilde{\theta}\right)$ costs $O\left(d^{2}\right)$
- Directional derivative $H_{i}(\tilde{\theta})\left(\theta_{t}-\tilde{\theta}\right) \operatorname{costs} O\left(\operatorname{eval}\left(f_{i}\right)\right)$

Build approximations $\hat{H}_{j}(\theta) \approx H_{j}(\theta)$

## Different ways to approximate the Hessian

$$
\hat{H}_{i}(\theta) \approx H_{i}(\theta)
$$



We tried:

- Diagonal approximations
- Rank-1 approximation based on secant equation
- Low rank approximations using Sketching and projecting


## Sketching the stochastic Hessian



Sketching matrix
$S \sim \mathcal{D}$ fixed distribution $S \in \mathbb{R}^{d \times \tau}$

Costs $\tau \times O\left(\operatorname{eval}\left(f_{i}\right)\right)$ to evaluate $H_{i}(\theta) S$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

find $X$ such that

$$
X S=H_{i} S
$$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

find $X$ such that

$$
X S=H_{i} S, \quad X=X^{\top}
$$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
& \hat{H}_{i}=\arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \quad \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
& \text { where }\|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta})
\end{aligned}
$$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \quad \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
\text { where } & \|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta}) \\
\hat{H}_{i}= & H S\left(S^{T} H S\right)^{-1} S^{\top} H_{i}\left(I-S\left(S^{T} H S\right)^{-1} S^{\top} H\right) \\
& +H_{i} S\left(S^{T} H S\right)^{-1} S^{\top} H .
\end{aligned}
$$

Total inner iteration costs: $O\left(\tau \times \operatorname{eval}\left(f_{i}\right)+\tau^{2} d+\tau^{3}\right)$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
\text { where } & \|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta}) \\
\hat{H}_{i}= & H S\left(S^{T} H S\right)^{-1} S^{\top} H_{i}\left(I-S\left(S^{T} H S\right)^{-1} S^{\top} H\right) \\
& +H_{i} S\left(S^{T} H S\right)^{-1} S^{\top} H . \quad \text { rank } 2 \tau
\end{aligned}
$$

Total inner iteration costs: $O\left(\tau \times \operatorname{eval}\left(f_{i}\right)+\tau^{2} d+\tau^{3}\right)$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
\text { where } & \|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta}) \\
\hat{H}_{i}= & H S\left(S^{T} H S\right)^{-1} S^{\top} H_{i}\left(I-S\left(S^{T} H S\right)^{-1} S^{\top} H\right) \\
& +H_{i} S\left(S^{T} H S\right)^{-1} S^{\top} H . \quad \text { rank } 2 \tau
\end{aligned}
$$

Total inner iteration costs: $O\left(\tau \times \operatorname{eval}\left(f_{i}\right)+\tau^{2} d+\tau^{3}\right)$

$$
\begin{aligned}
& \frac{1}{n} \sum_{j=1}^{n} \hat{H}_{j}=H S\left(S^{T} H S\right)^{-1} S^{\top} H \\
& \quad \text { Total outer costs: } O\left(n \tau \times \operatorname{eval}\left(f_{i}\right)\right)
\end{aligned}
$$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
\text { where } & \|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta}) \\
\hat{H}_{i}= & H S\left(S^{T} H S\right)^{-1} S^{\top} H_{i}\left(I-S\left(S^{T} H S\right)^{-1} S^{\top} H\right) \\
& +H_{i} S\left(S^{T} H S\right)^{-1} S^{\top} H . \quad \text { rank } 2 \tau
\end{aligned}
$$

Total inner iteration costs: $O\left(\tau \times \operatorname{eval}\left(f_{i}\right)+\tau^{2} d+\tau^{3}\right)$

$$
\begin{aligned}
& \frac{1}{n} \sum_{j=1}^{n} \hat{H}_{j}=H S\left(S^{T} H S\right)^{-1} S^{\top} H \\
& \quad \text { Total outer costs: } O\left(n \tau \times \operatorname{eval}\left(f_{i}\right)\right)
\end{aligned}
$$

## Sketching and Projecting the Hessian: Action Matching (AM) approximation

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \text { subject to } X S=H_{i} S, \quad X=X^{\top} \\
\text { where } & \|X\|_{F(H)}^{2} \stackrel{\text { def }}{=} \operatorname{Tr}\left(X H X^{\top} H\right) \text { and } H=\nabla^{2} f(\tilde{\theta}) \\
\hat{H}_{i}= & H S\left(S^{T} H S\right)^{-1} S^{\top} H_{i}\left(I-S\left(S^{T} H S\right)^{-1} S^{\top} H\right) \\
& +H_{i} S\left(S^{T} H S\right)^{-1} S^{\top} H . \quad \text { rank } 2 \tau
\end{aligned}
$$

Total inner iteration costs: $O\left(\tau \times \operatorname{eval}\left(f_{i}\right)+\tau^{2} d+\tau^{3}\right)$

$$
\begin{gathered}
\frac{1}{n} \sum_{j=1}^{n} \hat{H}_{j}=H S\left(S^{T} H S\right)^{-1} S^{\top} H . \quad \text { What about } S ? \\
\text { Total outer costs: } O\left(n \tau \times \operatorname{eval}\left(f_{i}\right)\right)
\end{gathered}
$$

## Choosing the sketch matrix

$$
\begin{aligned}
\hat{H}_{i}= & \arg \min _{X \in \mathbb{R}^{d \times d}}\|X\|_{F(H)}^{2} \\
& \text { subject to } \quad X S=H_{i} S, \quad X=X^{\top}
\end{aligned}
$$

AMgauss: $\quad S \sim \mathcal{N}(0, I)$ has Gaussian entries samplied i.i.d at each iteration

AMprev: Averages of previous search directions

$$
S=\left[\bar{g}_{0}, \ldots, \bar{g}_{\tau-1}\right] \quad \text { where } \quad \bar{g}_{i}=\frac{\tau}{m} \sum_{j=\frac{m}{\tau} i}^{\frac{m}{\tau}(i+1)-1} g_{j}
$$

## AM: Experiment works well



## AM: Experiment works ok



## AM: Experiment works badly



## Take home:

Can use Hessian to diminish variance

Speed-ups with less gain and risk compared to Newton type methods.

New compressed Hessian estimates using sketching and projecting

E gowerrobert/StochOpt.jl

## PDF

Bruce Christianson. Automatic Hessians by reverse accumulation. In: IMA Journal of
Numerical Analysis 12.2 (1992), pp. 135-150.

RMG and P. Richtárik, Randomized quasi-Newton updates are linearly convergent matrix inversion algorithms, SIAM Journal on Matrix Analysis and Applications, 38(4), 1380-1409, 2017
D.. Goldfarb, (1970). A Family of Variable-Metric Methods Derived by Variational Means.
Mathematics of Computation, 24(109), 23.

