
Supplementary Material for Dual Space Gradient Descent for Online Learning

Trung Le, Tu Dinh Nguyen, Vu Nguyen, Dinh Phung
Centre for Pattern Recognition and Data Analytics
Deakin University, Australia

{trung.l, tu.nguyen, v.nguyen, dinh.phung}@deakin.edu.au

1 Suitability of Loss Functions

In this section, we present the suitability of the loss functions for Hinge, smooth Hinge, and Logistic for classification and ℓ_1 , and ε -insensitive for regression. We prove that these losses satisfy the condition: there exists a positive constant A such that $|\nabla_o l(y, o)| \leq A, \forall y, o$. For each loss, we show its two forms used in the paper w.r.t o and \mathbf{w} .

Hinge loss

$$\begin{aligned} l(y, o) &= \max(0, 1 - yo) \\ l(\mathbf{w}, \mathbf{x}, y) &= \max(0, 1 - y\mathbf{w}^\top \Phi(\mathbf{x})) \\ \nabla_o l(y, o) &= -\mathbb{I}_{yo \leq 1} y \\ |\nabla_o l(y, o)| &= |\mathbb{I}_{yo \leq 1}| \leq 1 = A \end{aligned}$$

Logistic loss

$$\begin{aligned} l(y, o) &= \log(1 + e^{-yo}) \\ l(\mathbf{w}, \mathbf{x}, y) &= \log(1 + e^{-y\mathbf{w}^\top \Phi(\mathbf{x})}) \\ \nabla_o l(y, o) &= \frac{-ye^{-yo}}{e^{-yo} + 1} \\ |\nabla_o l(y, o)| &= \left| \frac{e^{-yo}}{e^{-yo} + 1} \right| < 1 = A \end{aligned}$$

Smooth Hinge loss [4]

$$\begin{aligned} l(y, o) &= \begin{cases} 0 & \text{if } yo > 1 \\ 1 - yo - \frac{\tau}{2} & \text{if } yo < 1 - \tau \\ \frac{1}{2\tau} (1 - yo)^2 & \text{otherwise} \end{cases} \\ l(\mathbf{w}, \mathbf{x}, y) &= \begin{cases} 0 & \text{if } y\mathbf{w}^\top \Phi(\mathbf{x}) > 1 \\ 1 - y\mathbf{w}^\top \Phi(\mathbf{x}) - \frac{\tau}{2} & \text{if } y\mathbf{w}^\top \Phi(\mathbf{x}) < 1 - \tau \\ \frac{1}{2\tau} (1 - y\mathbf{w}^\top \Phi(\mathbf{x}))^2 & \text{otherwise} \end{cases} \\ \nabla_o l(y, o) &= -\mathbb{I}_{\{yo < 1 - \tau\}} y + \tau^{-1} \mathbb{I}_{1 - \tau \leq yo \leq 1} (yo - 1) y \\ |\nabla_o l(y, o)| &= |\mathbb{I}_{\{yo < 1 - \tau\}}| + |\tau^{-1} \mathbb{I}_{1 - \tau \leq yo \leq 1} (yo - 1)| \\ &\leq |\mathbb{I}_{\{yo < 1 - \tau\}}| + \tau^{-1} \tau |\mathbb{I}_{1 - \tau \leq yo \leq 1}| \leq 1 = A \end{aligned}$$

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain

Acknowledgment: This work is partially supported by the Australian Research Council under the Discovery Project DP160109394.

ℓ_1 loss

$$\begin{aligned} l(y, o) &= |y - o| \\ l(\mathbf{w}, \mathbf{x}, y) &= |y - \mathbf{w}^\top \Phi(\mathbf{x})| \\ \nabla_o l(y, o) &= \mathbf{sign}(o - y) \\ |\nabla_o l(y, o)| &\leq 1 = A \end{aligned}$$

ε -insensitive loss

$$\begin{aligned} l(y, o) &= \max(0, |y - o| - \varepsilon) \\ l(\mathbf{w}, \mathbf{x}, y) &= \max(0, |y - \mathbf{w}^\top \Phi(\mathbf{x})| - \varepsilon) \\ \nabla_o l(y, o) &= \mathbb{I}_{|y-o| \geq \varepsilon} \mathbf{sign}(o - y) \\ |\nabla_o l(y, o)| &\leq 1 = A \end{aligned}$$

We note that \mathbb{I}_A denotes the indicator function which renders 1 if A is true and 0 otherwise.

2 Proofs

Lemma 1. *After the iteration t , we have the following representations*

$$\hat{\mathbf{w}}_t = \sum_{j=1}^t \alpha_j (1 - \beta_j) \Phi(\mathbf{x}_j) \quad (1)$$

$$\tilde{\mathbf{w}}_t = \sum_{j=1}^t \alpha_j \beta_j \mathbf{z}(\mathbf{x}_j) \quad (2)$$

$$\mathbf{w}_t = \sum_{j=1}^t \alpha_j \Phi(\mathbf{x}_j) \quad (3)$$

where $\alpha_j = -\eta_t \nabla_o l(y_j, f_j^h(\mathbf{x}_j))$, $\forall j = 1, \dots, t$ and $\eta_t = \frac{1}{\lambda t}$.

Proof. Since if $\beta_j = 1$, we perform the budget maintenance procedure and move the current vector to the random-feature space, we have the representations in Eqs. (1,2,3). In addition at the iteration j , $\Phi(\mathbf{x}_j)$ arrives with the initial coefficient $\alpha_j = -\eta_j \nabla_o l(y_j, f_j^h(\mathbf{x}_j))$. After the iteration $t > j$, this coefficient becomes

$$\alpha_j = -\frac{t-1}{t} \frac{t-2}{t-1} \dots \frac{j}{j+1} \frac{1}{\lambda j} \nabla_o l(y_i, f_j^h(\mathbf{x}_j)) = -\eta_t \nabla_o l(y_j, f_j^h(\mathbf{x}_j))$$

□

Theorem 2. *With a probability at least $1 - 2^8 \left(\frac{\sigma_\mu A d_\mathcal{X}}{\lambda \varepsilon}\right) \exp\left(-\frac{D \lambda^2 \varepsilon^2}{4(M+2)A^2}\right)$ where $d_\mathcal{X}$ specifies the diameter of the compact set \mathcal{X} , we have*

i) $|f_t(\mathbf{x}) - f_t^h(\mathbf{x})| \leq \varepsilon$ for all $t > 0$ and $\mathbf{x} \in \mathcal{X}$.

ii) $\mathbb{E}[|f_t(\mathbf{x}) - f_t^h(\mathbf{x})|] \leq A^{-1} \lambda \varepsilon \sum_{j=1}^t \mathbb{E}[\alpha_j^2]^{1/2} \mu_j^{1/2}$ where $\mu_j = p(\beta_j = 1)$.

Let us define a random map $z : \mathbb{R}^d \rightarrow \mathbb{R}^{2D}$ where $z(x) = \frac{1}{D^{1/2}} [\cos(\omega_i^\top x), \sin(\omega_i^\top x)]_{i=1}^D$ and $\omega_1, \dots, \omega_D \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \sigma^{-2}I)$ for every $x \in \mathbb{R}^d$. We would like to restate Claim 1 in [3].

Let \mathcal{M} be a compact subset of \mathbb{R}^d with diameter $\text{diam}(\mathcal{M})$. Then, for the random mapping $z(\cdot)$, we have

$$\mathbb{P}\left(\sup_{x, x' \in \mathcal{M}} |K(x, x') - z(x)^\top z(x')| < \varepsilon\right) \geq 1 - 2^8 \left(\frac{\sigma \text{diam}(\mathcal{M})}{\varepsilon}\right) \exp\left(\frac{-D\varepsilon^2}{4(d+2)}\right)$$

where $K(x, x') = e^{-\frac{\|x-x'\|^2}{2\sigma^2}}$.

Proof. We denote

$$\omega = (\omega_1, \dots, \omega_D) \sim p_\omega(\omega) = \prod_{i=1}^D \mathcal{N}(\omega_i | 0, \sigma^{-2}I)$$

$$\tilde{K}(x, x') = z(x)^\top z(x') = D^{-1} \sum_{i=1}^D \left(\cos(\omega_i^\top x) \cos(\omega_i^\top x') + \sin(\omega_i^\top x) \sin(\omega_i^\top x') \right)$$

We further denote

$$g(\omega) = \sup_{x, x' \in \mathcal{M}} \left| K(x, x') - \tilde{K}(x, x') \right|$$

$$G_\varepsilon = \{ \omega : g(\omega) < A^{-1}\lambda\varepsilon \}$$

It is certain that $\mathbb{P}_\omega(G_\varepsilon) \geq 1 - \theta$ where $\theta = 2^8 \left(\frac{\sigma \text{Adiam}(\mathcal{M})}{\lambda\varepsilon} \right) \exp\left(\frac{-D\lambda^2\varepsilon^2}{4(d+2)A^2} \right)$ and for every $\omega \in G_\varepsilon$ and $x, x' \in \mathcal{M}$ we have

$$\left| K(x, x') - \tilde{K}(x, x') \right| < A^{-1}\lambda\varepsilon$$

We now turn back to Theorem 2. It appears that

$$\left| f_t(x) - f_t^h(x) \right| \leq \sum_{j=1}^t \beta_j |\alpha_j| \left| K(x_j, x) - \tilde{K}(x_j, x) \right|$$

Therefore, for every $\omega \in G_\varepsilon$ we have

$$\left| f_t(x) - f_t^h(x) \right| \leq A^{-1}\lambda\varepsilon \sum_{j=1}^t \beta_j |\alpha_j|$$

Let us denote $s = (x_1, y_1), \dots, (x_t, y_t)$. Taking expectation of the above inequality w.r.t s , we gain for all $\omega \in G_\varepsilon$

$$\begin{aligned} \mathbb{E}_s \left[\left| f_t(x) - f_t^h(x) \right| \right] &\leq A^{-1}\lambda\varepsilon \sum_{j=1}^t \mathbb{E}_s [\beta_j^2]^{1/2} \mathbb{E}_s [\alpha_j^2]^{1/2} \\ &\leq A^{-1}\lambda\varepsilon \sum_{j=1}^t \mu_j \mathbb{E}_s [\alpha_j^2]^{1/2} \end{aligned}$$

It means that

$$\mathbb{P}_\omega \left(\mathbb{E}_s \left[\left| f_t(x) - f_t^h(x) \right| \right] \leq A^{-1}\lambda\varepsilon \sum_{j=1}^t \mu_j \mathbb{E}_s [\alpha_j^2]^{1/2} \right) \geq \mathbb{P}_\omega(G_\varepsilon) \geq 1 - \theta$$

□

Lemma 3. *The following statement holds for all t*

$$\|\mathbf{w}_t\| \leq \frac{A}{\lambda}$$

Proof. Using Lemma 1, we have

$$\mathbf{w}_t = \sum_{j=1}^t \alpha_j \Phi(\mathbf{x}_j)$$

where $\alpha_j = -\eta_t \nabla_{o} l(y_j, f_j^h(\mathbf{x}_j))$.

It implies that

$$\|\mathbf{w}_t\| \leq \sum_{j=1}^t |\alpha_j| \|\Phi(\mathbf{x}_j)\| \leq \sum_{j=1}^t |\alpha_j| \leq \sum_{j=1}^t \frac{A}{\lambda t} = \frac{A}{\lambda}$$

□

Lemma 4. *The following statement holds for all t*

$$\|g_t\| \leq G = 2A$$

where we define $g_t = \lambda \mathbf{w}_t + \nabla_{\mathbf{w}} l(\mathbf{w}_t, \mathbf{x}_t, y_t) = \lambda \mathbf{w}_t + \nabla_{o} l(y_t, f_t(\mathbf{x}_t)) \Phi(\mathbf{x}_t)$.

Proof. We derive as

$$\|g_t\| \leq \lambda \|\mathbf{w}_t\| + \|\nabla_{o} l(y_t, f_t(\mathbf{x}_t)) \Phi(\mathbf{x}_t)\| \leq \lambda \frac{A}{\lambda} + A = 2A$$

□

Lemma 5. *The following statement holds for all t*

$$\mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] \leq W^2$$

where $W = \frac{2A(1+\sqrt{5})}{\lambda}$.

Proof. Recall that $g_t = \lambda \mathbf{w}_t + \nabla_{\mathbf{w}} l(\mathbf{w}_t, \mathbf{x}_t, y_t) = \lambda \mathbf{w}_t + \nabla_{o} l(y_t, f_t(\mathbf{x}_t)) \Phi(\mathbf{x}_t)$. It is obvious that g_t satisfies

$$\mathbb{E}_{(\mathbf{x}_t, y_t)} [g_t | \mathbf{w}_t] = \mathcal{J}'(\mathbf{w}_t)$$

We have the following if we denote $\delta g_t = g_t - g_t^h$

$$\begin{aligned} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &= \|\mathbf{w}_t - \eta_t g_t^h - \mathbf{w}^*\|^2 = \|\mathbf{w}_t - \eta_t g_t - \mathbf{w}^* + \eta_t \delta g_t\|^2 \\ &= \|\mathbf{w}_t - \mathbf{w}^*\|^2 - 2\eta_t g_t^\top (\mathbf{w}_t - \mathbf{w}^*) + \eta_t^2 \|g_t\|^2 - 2\eta_t^2 g_t^\top \delta g_t + \eta_t^2 \|\delta g_t\|^2 + 2\eta_t (\mathbf{w}_t - \mathbf{w}^*)^\top \delta g_t \end{aligned}$$

It appears that

$$\begin{aligned} \delta g_t &= [\nabla_{o} l(y_t, f_t(\mathbf{x}_t)) - \nabla_{o} l(y_t, f_t^h(\mathbf{x}_t))] \Phi(\mathbf{x}_t) \\ \|\delta g_t\| &= |\nabla_{o} l(y_t, f_t(\mathbf{x}_t)) - \nabla_{o} l(y_t, f_t^h(\mathbf{x}_t))| \leq 2A \end{aligned}$$

Hence, we obtain

$$\begin{aligned} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &\leq \|\mathbf{w}_t - \mathbf{w}^*\|^2 - 2\eta_t g_t^\top (\mathbf{w}_t - \mathbf{w}^*) + \eta_t^2 G^2 + 4\eta_t^2 GA + 4\eta_t^2 A^2 \\ &\quad + 2\eta_t \|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\| \end{aligned}$$

Taking conditional expectation w.r.t \mathbf{w}_t on both sides of the above inequality, we gain

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \right] &\leq \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] - 2\eta_t \nabla_{\mathbf{w}} \mathcal{J}(\mathbf{w}_t)^\top (\mathbf{w}_t - \mathbf{w}^*) + \eta_t^2 G^2 + 4\eta_t^2 GA \\ &\quad + 4\eta_t^2 A^2 + 2\eta_t \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|] \\ &\leq \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] + 16A^2 \eta_t^2 + 2\eta_t \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|] - \frac{1}{t} \|\mathbf{w}_t - \mathbf{w}^*\|^2 \end{aligned}$$

Here we note that we have used

$$\nabla_{\mathbf{w}} \mathcal{J}(\mathbf{w}_t)^\top (\mathbf{w}_t - \mathbf{w}^*) \geq \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) + \frac{\lambda}{2} \|\mathbf{w}_t - \mathbf{w}^*\|^2 \geq \frac{\lambda}{2} \|\mathbf{w}_t - \mathbf{w}^*\|^2$$

Taking expectation on both sides again, we obtain

$$\begin{aligned}\mathbb{E} \left[\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \right] &\leq \frac{t-1}{t} \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] + \frac{16A^2}{\lambda^2 t^2} + \frac{4A \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right]^{1/2}}{\lambda t} \\ &\leq \frac{t-1}{t} \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] + \frac{16A^2}{\lambda^2 t} + \frac{4A \mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right]^{1/2}}{\lambda t}\end{aligned}$$

Choose $W = \frac{2A(1+\sqrt{5})}{\lambda}$, we have if $\mathbb{E} \left[\|\mathbf{w}_t - \mathbf{w}^*\|^2 \right] \leq W^2$ then $\mathbb{E} \left[\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \right] \leq W^2$. \square

Theorem 6. *The following statement guarantees for all T*

$$\mathbb{E} [\mathcal{J}(\bar{\mathbf{w}}_T) - \mathcal{J}(\mathbf{w}^*)] \leq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) \right] \leq \frac{8A^2 (\log(T) + 1)}{\lambda T} + \frac{1}{T} W \sum_{t=1}^T \mathbb{E} [M_t^2]^{1/2}$$

where $\bar{\mathbf{w}}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{w}_t$, $M_t = \nabla_o l(y_t, f_t(\mathbf{x}_t)) - \nabla_o l(y_t, f_t^h(\mathbf{x}_t))$.

Proof. Recall that $g_t = \lambda \mathbf{w}_t + \nabla_{\mathbf{w}} l(\mathbf{w}_t, \mathbf{x}_t, y_t) = \lambda \mathbf{w}_t + \nabla_o l(y_t, f_t(\mathbf{x}_t)) \Phi(\mathbf{x}_t)$. It is obvious that g_t satisfies

$$\mathbb{E}_{(\mathbf{x}_t, y_t)} [g_t | \mathbf{w}_t] = \nabla_{\mathbf{w}} \mathcal{J}(\mathbf{w}_t)$$

We have the following if we denote $\delta g_t = g_t - g_t^h$

$$\begin{aligned}\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &= \|\mathbf{w}_t - \eta_t g_t^h - \mathbf{w}^*\|^2 = \|\mathbf{w}_t - \eta_t g_t - \mathbf{w}^* + \eta_t \delta g_t\|^2 \\ &= \|\mathbf{w}_t - \mathbf{w}^*\|^2 - 2\eta_t g_t^\top (\mathbf{w}_t - \mathbf{w}^*) + \eta_t^2 \|g_t\|^2 - 2\eta_t^2 g_t^\top \delta g_t + \eta_t^2 \|\delta g_t\|^2 + 2\eta_t (\mathbf{w}_t - \mathbf{w}^*)^\top \delta g_t\end{aligned}$$

It appears that

$$\begin{aligned}\delta g_t &= [\nabla_o l(y_t, f_t(\mathbf{x}_t)) - \nabla_o l(y_t, f_t^h(\mathbf{x}_t))] \Phi(\mathbf{x}_t) \\ \|\delta g_t\| &= |\nabla_o l(y_t, f_t(\mathbf{x}_t)) - \nabla_o l(y_t, f_t^h(\mathbf{x}_t))| \leq 2A\end{aligned}$$

Hence, we obtain

$$\begin{aligned}\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &\leq \|\mathbf{w}_t - \mathbf{w}^*\|^2 - 2\eta_t g_t^\top (\mathbf{w}_t - \mathbf{w}^*) + \eta_t^2 G^2 + 4\eta_t^2 GA + 4\eta_t^2 A^2 \\ &\quad + 2\eta_t \|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\| \\ g_t^\top (\mathbf{w}_t - \mathbf{w}^*) &\leq \frac{\|\mathbf{w}_t - \mathbf{w}^*\|^2 - \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2}{2\eta_t} + 8A^2 \eta_t + \|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|\end{aligned}$$

Taking conditional expectation w.r.t \mathbf{w}_t on both sides, we gain

$$\begin{aligned}\nabla_{\mathbf{w}} \mathcal{J}(\mathbf{w}_t)^\top (\mathbf{w}_t - \mathbf{w}^*) &\leq \mathbb{E} \left[\frac{\|\mathbf{w}_t - \mathbf{w}^*\|^2}{2\eta_t} \right] - \mathbb{E} \left[\frac{\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2}{2\eta_t} \right] + 8A^2 \eta_t + \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|] \\ \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) + \frac{\lambda}{2} \|\mathbf{w}_t - \mathbf{w}^*\|^2 &\leq \mathbb{E} \left[\frac{\|\mathbf{w}_t - \mathbf{w}^*\|^2}{2\eta_t} \right] - \mathbb{E} \left[\frac{\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2}{2\eta_t} \right] \\ &\quad + 8A^2 \eta_t + \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|]\end{aligned}$$

Taking expectation on both sides once again, we achieve

$$\begin{aligned}\mathbb{E} [\mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*)] &\leq \frac{\lambda}{2} (t-1) \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2] - \frac{\lambda}{2} t \mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2] \\ &\quad + 8A^2 \eta_t + \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\| \|\delta g_t\|] \\ \mathbb{E} [\mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*)] &\leq \frac{\lambda}{2} (t-1) \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2] - \frac{\lambda}{2} t \mathbb{E} [\|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2] \\ &\quad + 8A^2 \eta_t + \mathbb{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2]^{1/2} \mathbb{E} [\|\delta g_t\|^2]^{1/2}\end{aligned}$$

Taking sum the above inequality when $t = 1, \dots, T$, we obtain

$$\begin{aligned} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) \right] &\leq \frac{8A^2}{\lambda} \sum_{t=1}^T \frac{1}{t} + \frac{1}{T} W \sum_{t=1}^T \mathbb{E} [M_t^2]^{1/2} \\ &\leq \frac{8A^2 (\log T + 1)}{\lambda T} + \frac{1}{T} W \sum_{t=1}^T \mathbb{E} [M_t^2]^{1/2} \end{aligned}$$

Here we note that

$$\|\delta g_t\| = \left\| \left[\nabla_{ol}(y_t, f_t(\mathbf{x}_t)) - \nabla_{ol}(y_t, f_t^h(\mathbf{x}_t)) \right] \Phi(\mathbf{x}_t) \right\| = |M_t|$$

The last conclusion comes from the convexity of the function $\mathcal{J}(\cdot)$. \square

Theorem 7. Assume that $l(y, o)$ is a γ -strongly smooth loss function. With a probability at least $1 - \theta$, the following statements hold

$$i) \quad \mathbb{E} [\mathcal{J}(\bar{\mathbf{w}}_T) - \mathcal{J}(\mathbf{w}^*)] \leq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) \right] \leq \frac{8A^2 (\log T + 1)}{\lambda T} + \frac{1}{T} W \gamma \varepsilon \sum_{t=1}^T \left(\frac{\sum_{i=1}^t \mu_i}{t} \right)^{1/2}$$

$$ii) \quad \mathbb{E} [\mathcal{J}(\bar{\mathbf{w}}_T) - \mathcal{J}(\mathbf{w}^*)] \leq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) \right] \leq \frac{8A^2 (\log T + 1)}{\lambda T} + W \gamma \varepsilon$$

$$\text{where } \theta = 2^8 \left(\frac{\sigma_\mu A d_X}{\lambda \varepsilon} \right) \exp \left(-\frac{D \lambda^2 \varepsilon^2}{4(M+2)A^2} \right).$$

Proof. From the smoothness of the loss function, we have

$$\left| \nabla_{ol}(y_t, f_t(\mathbf{x}_t)) - \nabla_{ol}(y_t, f_t^h(\mathbf{x}_t)) \right| \leq \gamma |f_t(\mathbf{x}_t) - f_t^h(\mathbf{x}_t)|$$

Referring to Lemma 2, with a probability at least $1 - 2^8 \left(\frac{\sigma_\mu A d_X}{\lambda \varepsilon} \right) \exp \left(-\frac{D \lambda^2 \varepsilon^2}{4(M+2)A^2} \right) = 1 - \theta$ we have

$$\begin{aligned} |M_t| &\leq \gamma A^{-1} \lambda \varepsilon \sum_{j=1}^t |\alpha_j| \beta_j \leq \gamma A^{-1} \lambda \varepsilon \sum_{j=1}^t \frac{A}{\lambda t} \beta_j = \frac{\gamma \varepsilon}{t} \sum_{j=1}^t \beta_j \\ M_t^2 &\leq \frac{\gamma^2 \varepsilon^2}{t^2} \left(\sum_{j=1}^t \beta_j \right)^2 \leq \frac{\gamma^2 \varepsilon^2}{t} \sum_{j=1}^t \beta_j^2 = \frac{\gamma^2 \varepsilon^2}{t} \sum_{j=1}^t \beta_j \quad (\text{since } \beta_i = 0 \text{ or } 1) \\ \mathbb{E} [M_t^2] &\leq \frac{\gamma^2 \varepsilon^2}{t} \left(\sum_{j=1}^t \mu_j \right) \end{aligned}$$

and $|M_t| \leq \gamma \varepsilon$. Therefore, with a probability at least $1 - \theta$ we achieve

$$\begin{aligned} \mathbb{E} [\mathcal{J}(\bar{\mathbf{w}}_T) - \mathcal{J}(\mathbf{w}^*)] &\leq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*) \right] \\ &\leq \frac{8A^2 (\log T + 1)}{\lambda T} + \frac{1}{T} W \gamma \varepsilon \sum_{t=1}^T \frac{\left(\sum_{j=1}^t \mu_j \right)^{1/2}}{t^{1/2}} \\ &\leq \frac{8A^2 (\log T + 1)}{\lambda T} + \frac{1}{T} W \gamma \varepsilon \sum_{t=1}^T \left(\frac{\sum_{j=1}^t \mu_j}{t} \right)^{1/2} \end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}[\mathcal{J}(\bar{\mathbf{w}}_T) - \mathcal{J}(\mathbf{w}^*)] &\leq \mathbb{E}\left[\frac{1}{T} \sum_{t=1}^T \mathcal{J}(\mathbf{w}_t) - \mathcal{J}(\mathbf{w}^*)\right] \\
&\leq \frac{8A^2(\log T + 1)}{\lambda T} + \frac{1}{T} W \sum_{t=1}^T \gamma \varepsilon \\
&\leq \frac{8A^2(\log T + 1)}{\lambda T} + W\gamma\varepsilon
\end{aligned}$$

□

3 Computational Complexities of DualSGD and FOGD

We compare the computational complexities of our proposed DualSGD and Fourier Online Gradient Descent (FOGD) [2]. Recall that M and D denote the dimensions of input space and feature space, and B the budget size. There are four operators: (i) random feature mapping; (ii) kernel function; (iii) sorting coefficients of support vectors and (iv) prediction. The random feature mapping first projects the input data vector to random feature space with $\mathcal{O}(MD)$ computational complexity, and then compute \sin, \cos on the random feature dimension with $\mathcal{O}(2D * 2^{\log n} n \log^2 n)$ where n is the number of bits accuracy [1]. The kernel function, sorting coefficients and prediction operate in $\mathcal{O}(MB)$, $\mathcal{O}(B \log B)$ and $\mathcal{O}(D)$ complexity, respectively. The FOGD performs random feature mapping and prediction whilst the DualSGD performs all four operators.

Let D_1 and D_2 denote the number of random features of FOGD and DualSGD. The computational complexities of FOGD and DualSGD reads

$$\begin{aligned}
\mathcal{O}_{\text{FOGD}} &= \mathcal{O}(MD_1 + 2D_1 * 2^{\log n} n \log^2 n + D_1) = U(MD_1 + 2D_1 * 2^{\log n} n \log^2 n + D_1) \\
\mathcal{O}_{\text{DualSGD}} &= \mathcal{O}(MD_2 + 2D_2 * 2^{\log n} n \log^2 n + D_2 + MB + B \log B) \\
&= V(MD_2 + 2D_2 * 2^{\log n} n \log^2 n + D_2 + MB + B \log B)
\end{aligned}$$

where U, V are the number of iterations.

Taking the subtraction of $\mathcal{O}_{\text{FOGD}}$ and $\mathcal{O}_{\text{DualSGD}}$, we obtain:

$$\begin{aligned}
\hat{\mathcal{O}} &= \mathcal{O}_{\text{FOGD}} - \mathcal{O}_{\text{DualSGD}} \\
&= M(UD_1 - VD_2 - B) + (UD_1 - VD_2)(2 * 2^{\log n} n \log^2 n + 1) - B \log B
\end{aligned}$$

According Fig. 1 in the introduction section, $D_1 \gg D_2$ and $D_1 \gg B$, thus $D_1 - D_2 \gg B$. In addition, we assume that $U = V$ and normally use double-precision floating-point with $n = 64$ (bits) for storing and computing real number, thus $2 * 2^{\log n} n \log^2 n + 1 > \log B$. Finally, we can see that $\hat{\mathcal{O}} \gg 0$, thus the computational complexity of DualSGD, in practice, is significantly lower than that of FOGD.

References

- [1] R. P Brent and P. Zimmermann. *Modern computer arithmetic*, volume 18. Cambridge University Press, 2010.
- [2] J. Lu, S. C.H. Hoi, J. Wang, P. Zhao, and Z.-Y. Liu. Large scale online kernel learning. *J. Mach. Learn. Res.*, 2015.
- [3] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.
- [4] S. Shalev-Shwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss. *Journal of Machine Learning Research*, 14(1):567–599, 2013.