

FFTLasso: Large-Scale LASSO in the Fourier Domain (Supplementary Material)

Adel Bibi, Hani Itani, Bernard Ghanem

King Abdullah University of Science and Technology (KAUST), Saudi Arabia

adel.bibi@kaust.edu.sa, hmi10@mail.aub.edu, bernard.ghanem@kaust.edu.sa

In this supplementary, we will discuss the details to the provided formulation in the main manuscript. Next, we will present an efficient matrix notation implementation of FFT-Lasso. Lastly, we discuss the memory efficiency of the algorithm. The original LASSO is given as follows:

$$\min_{\mathbf{c}} \|\mathbf{A}\mathbf{c} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{c}\|_1 \quad (1)$$

Problem 1 is equivalent to the FFTLasso reformulation:

$$\min_{\tilde{\mathbf{c}}} \|\tilde{\mathbf{A}}\tilde{\mathbf{c}} - \mathbf{b}\|_2^2 + \lambda \|\tilde{\mathbf{c}}\|_1 \quad \text{s.t.} \quad \mathbf{D}\tilde{\mathbf{c}} = \mathbf{0}, \quad (2)$$

where $\tilde{\mathbf{A}} \in \mathbb{R}^{m \times mn}$ and is block wise circulant. First and for reasons discussed in the main manuscript, we need to find the dual equivalence to problem 2. To do so, we introduce a new variable $\mathbf{r} = \tilde{\mathbf{A}}\tilde{\mathbf{c}} - \mathbf{b}$ as follows:

$$\min_{\tilde{\mathbf{c}}, \mathbf{r}} \|\mathbf{r}\|_2^2 + \lambda \|\tilde{\mathbf{c}}\|_1 \quad \text{s.t.} \quad \mathbf{D}\tilde{\mathbf{c}} = \mathbf{0}, \quad \tilde{\mathbf{A}}\tilde{\mathbf{c}} - \mathbf{b} = \mathbf{r} \quad (3)$$

Then, the corresponding lagrangian is defined as follows:

$$\begin{aligned} \mathcal{L}(\mathbf{r}, \tilde{\mathbf{c}}, \Psi, \theta) := & \|\mathbf{r}\|_2^2 + \lambda \|\tilde{\mathbf{c}}\|_1 \quad (4) \\ & + \Psi^T (\tilde{\mathbf{A}}\tilde{\mathbf{c}} - \mathbf{b} - \mathbf{r}) + \theta^T \mathbf{D}\tilde{\mathbf{c}} \end{aligned}$$

To find the dual problem, we minimize over the primal variables as follows:

$$\begin{aligned} \min_{\mathbf{r}} & \left[\|\mathbf{r}\|_2^2 - \Psi^T \mathbf{r} \right] - \Psi^T \mathbf{b} + \\ \min_{\tilde{\mathbf{c}}} & \left[\lambda \|\tilde{\mathbf{c}}\|_1 + \Psi^T \tilde{\mathbf{A}}\tilde{\mathbf{c}} + \theta^T \mathbf{D}\tilde{\mathbf{c}} \right] \quad (5) \end{aligned}$$

$$\begin{aligned} = & -\Psi^T \mathbf{b} - \frac{1}{4} \|\Psi\|_2^2 \\ & - \lambda \max_{\tilde{\mathbf{c}}} \left(\frac{\tilde{\mathbf{c}}^T (-\tilde{\mathbf{A}}^T \Psi - \mathbf{D}^T \theta) - \|\tilde{\mathbf{c}}\|_1}{\lambda} \right) \quad (6) \end{aligned}$$

$$= -\Psi^T \mathbf{b} - \frac{1}{4} \|\Psi\|_2^2 - \mathbb{1}_{\frac{1}{\lambda} \|\tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta\|_\infty \leq 1} \quad (7)$$

Therefore, the dual problem is given as follows:

$$\min_{\Psi, \theta} \Psi^T \mathbf{b} + \frac{1}{4} \|\Psi\|_2^2 \quad \text{s.t.} \quad \|\tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta\|_\infty \leq \lambda \quad (8)$$

Now, we apply ADMM to problem 8. First, we introduce the change of variables $\zeta = \tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta$ and append the $\|\cdot\|_\infty$ constraint as an indicator function. Problem 8 is then equivalent to:

$$\begin{aligned} \min_{\Psi, \theta, \zeta} & \frac{1}{4} \|\Psi\|_2^2 + \Psi^T \mathbf{b} + \mathbb{1}_{\|\zeta\|_\infty \leq \lambda} \\ \text{s.t.} & \zeta = \tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta \quad (9) \end{aligned}$$

The augmented lagrangian is defined as follows:

$$\begin{aligned} \mathcal{L}(\Psi, \theta, \zeta, \tilde{\mathbf{c}}) := & \frac{1}{4} \|\Psi\|_2^2 + \Psi^T \mathbf{b} + \mathbb{1}_{\|\zeta\|_\infty \leq \lambda} + \\ & \tilde{\mathbf{c}}^T (\tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta - \zeta) + \frac{\rho}{2} \|\tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta - \zeta\|_2^2 \quad (10) \end{aligned}$$

The primal-dual update variables are given as follows:
Update Ψ :

$$\begin{aligned} \Psi^{k+1} = & \operatorname{argmin}_{\Psi} \frac{1}{4} \|\Psi\|_2^2 + \Psi^T \mathbf{b} + \tilde{\mathbf{c}}^T \tilde{\mathbf{A}}^T \Psi \\ & + \frac{\rho}{2} \|\tilde{\mathbf{A}}^T \Psi + \mathbf{D}^T \theta - \zeta\|_2^2 \quad (11) \end{aligned}$$

The previous objective minimization reduces to solving the following linear system:

$$(\rho \tilde{\mathbf{A}} \tilde{\mathbf{A}}^T + \frac{1}{2} \mathbf{I}_m) \Psi^{k+1} = \tilde{\mathbf{A}} (\rho \zeta - \rho \mathbf{D}^T \theta - \tilde{\mathbf{c}}) - \mathbf{b} \quad (12)$$

Since $\tilde{\mathbf{A}} \mathbf{e} = \sum_i^N \mathbf{A}_i \mathbf{e}_i = \mathbf{F} \sum_i^N \hat{\mathbf{a}}_i^* \odot \hat{\mathbf{e}}_i^*$ where $\mathbf{e} = \rho \zeta - \rho \mathbf{D}^T \theta - \tilde{\mathbf{c}}$. The solution to the linear system is given as follows:

$$\mathbf{F} (\rho \sum_i^N \operatorname{diag}(\hat{\mathbf{a}}_i \odot \hat{\mathbf{a}}_i^*) + \frac{1}{2} \mathbf{I}_m) \mathbf{F}^H \Psi = \mathbf{F} \sum_i^N \operatorname{diag}(\hat{\mathbf{a}}_i^* \odot \hat{\mathbf{e}}_i^*) - \mathbf{b}$$

By multiplying both sides by \mathbf{F}^H where $\mathbf{F} \mathbf{F}^H = \mathbf{I}_m$, then:

$$(\rho \sum_i^N \operatorname{diag}(\hat{\mathbf{a}}_i \odot \hat{\mathbf{a}}_i^*) + \frac{1}{2} \mathbf{I}_m) \hat{\Psi}^* = \sum_i^N \operatorname{diag}(\hat{\mathbf{a}}_i^* \odot \hat{\mathbf{e}}_i^*) - \hat{\mathbf{b}}^*$$

$$\hat{\Psi}^* = \frac{\sum_i^N \hat{\mathbf{a}}_i^* \odot \hat{\mathbf{e}}_i^* - \hat{\mathbf{b}}^*}{\rho \sum_i^N \hat{\mathbf{a}}_i \odot \hat{\mathbf{a}}_i^* + \frac{1}{2}} \quad (13)$$

Update $\mathbf{D}^H\theta$:

$$\begin{aligned}\mathbf{D}^H\theta^{k+1} &= \operatorname{argmin}_{\mathbf{D}^H\theta} \tilde{\mathbf{c}}^H\mathbf{D}^H\theta + \frac{\rho}{2}\|\tilde{\mathbf{A}}^H\Psi + \mathbf{D}^H\theta - \zeta\|_2^2 \\ \mathbf{D}^H\theta^{k+1} &= \zeta - \tilde{\mathbf{c}}/\rho - \tilde{\mathbf{A}}^H\Psi\end{aligned}\quad (14)$$

To compute $\tilde{\mathbf{A}}^H\Psi$ efficiently, the diagonalization trick can be applied again as follows:

$$\tilde{\mathbf{A}}^H\Psi = \begin{bmatrix} \mathbf{F}(\hat{\mathbf{a}}_1 \odot \Psi^*) \\ \vdots \\ \mathbf{F}(\hat{\mathbf{a}}_n \odot \Psi^*) \end{bmatrix}\quad (15)$$

Update ζ :

$$\begin{aligned}\zeta^{k+1} &= \operatorname{argmin}_{\zeta} \mathbb{1}_{\|\zeta\|_\infty \leq \lambda} - \tilde{\mathbf{c}}^H\zeta + \frac{\rho}{2}\|\tilde{\mathbf{A}}^T\Psi + \mathbf{D}^T\theta - \zeta\|_2^2 \\ &= \operatorname{argmin}_{\zeta} \|\zeta - (\tilde{\mathbf{A}}^H\Psi + \mathbf{D}^T\theta + \tilde{\mathbf{c}}/\rho)\|_2^2 \\ &\text{s.t. } \|\zeta\|_\infty \leq \lambda\end{aligned}\quad (16)$$

The former is just a simple Euclidian projection to the $\|\cdot\|_\infty$ ball and is given as follows:

$$\zeta^{k+1} = \operatorname{sign}(\mathbf{t}) \odot \min(|\mathbf{t}|, \lambda)\quad (17)$$

where $\mathbf{t} = \tilde{\mathbf{A}}^T\Psi + \mathbf{D}^T\theta + \tilde{\mathbf{c}}/\rho$. Lastly, a Nesterov dual ascent is applied to the dual variables as follows:

Update $\tilde{\mathbf{c}}$:

$$\tilde{\mathbf{c}}_{k+1} = \tilde{\mathbf{y}}_k + \rho(\tilde{\mathbf{A}}^H\Psi + \mathbf{D}^H\theta - \zeta)\quad (18)$$

$$\tilde{\mathbf{y}}_{k+1} = (1+q)\tilde{\mathbf{c}}_{k+1} - q\tilde{\mathbf{c}}_k\quad (19)$$

The algorithm is given as follows:

Algorithm 1: FFTLasso for Problem (9)

Input : $\mathbf{b}, \mathbf{A}, \tilde{\mathbf{c}}_1 = \tilde{\mathbf{y}}_1 = \mathbf{e}_1 = \mathbf{t}_1 = \zeta_1 = \mathbf{D}^H\theta_1 = \mathbf{0}_{mn}, \Psi = \mathbf{0}_m, \lambda, \rho_1, \gamma > 1, q$.

Output: \mathbf{c}

while not converged **do**

compute: $\mathbf{e}_{k+1} = \rho_k\zeta_k - \rho_k\mathbf{D}^H\theta_k - \tilde{\mathbf{c}}_k$

$\hat{\Psi}^*$ update: $\hat{\Psi}_{k+1}^* = \frac{\sum_i^N \hat{\mathbf{a}}_i^* \odot \hat{\mathbf{e}}_{ik+1}^* - \hat{\mathbf{b}}^*}{\rho_k \sum_i^N \hat{\mathbf{a}}_i \odot \hat{\mathbf{a}}_i^* + \frac{1}{2}}$

compute: $\tilde{\mathbf{A}}^H\Psi_{k+1}$, see Eq (15)

$\mathbf{D}^H\theta$ update:

$(\mathbf{D}^H\theta_{k+1}) = (\zeta_k - \frac{1}{\rho_k}\tilde{\mathbf{c}}_k - \tilde{\mathbf{A}}^H\Psi_{k+1})$

$(\mathbf{D}^H\theta_{k+1})_{1:m:\text{end}} = \mathbf{0}_n$

compute: $\mathbf{t}_{k+1} = \tilde{\mathbf{A}}^H\Psi_{k+1} + \mathbf{D}^H\theta_{k+1} + \tilde{\mathbf{c}}_k/\rho_k$

ζ update: $\zeta_{k+1} = \operatorname{sign}(\mathbf{t}_{k+1}) \odot \min(|\mathbf{t}_{k+1}|, \lambda)$

$\tilde{\mathbf{c}}$ update:

$\tilde{\mathbf{c}}_{k+1} = \tilde{\mathbf{y}}_k + \rho_k(\tilde{\mathbf{A}}^H\Psi_{k+1} + \mathbf{D}^H\theta_{k+1} - \zeta_{k+1})$

compute: $\tilde{\mathbf{y}}_{k+1} = (1+q)\tilde{\mathbf{c}}_{k+1} - q\tilde{\mathbf{c}}_k$

$\rho_{k+1} = \gamma\rho_k$

end

$\mathbf{c} \leftarrow \tilde{\mathbf{c}}(1:m:\text{end})$

Next we present a more efficient implementation to FFT-Lasso algorithm using matrix notation. First we define the new variables in matrix notation as follows: $\mathbf{e} \in \mathbb{R}^{mn}$ is changed to $\mathbf{E} \in \mathbb{R}^{m \times n}$, $\mathbf{D}^H\theta$ is changed to $\Theta \in \mathbb{R}^{m \times n}$, $\mathbf{t} \in \mathbb{R}^{mn}$ is changed to $\mathbf{T} \in \mathbb{R}^{m \times n}$, $\mathbf{c} \in \mathbb{R}^{mn}$ is changed to $\mathbf{C} \in \mathbb{R}^{m \times n}$, $\mathbf{y} \in \mathbb{R}^{mn}$ is changed to $\mathbf{Y} \in \mathbb{R}^{m \times n}$, and lastly $\zeta \in \mathbb{R}^{mn}$ is changed to $\mathbf{Z} \in \mathbb{R}^{m \times n}$. The algorithm is then given as follows: In this matrix notation, there is no need

Algorithm 2: FFTLasso for Problem (9)

Input : $\mathbf{b}, \mathbf{A}, \hat{\mathbf{A}}^* = \mathbf{F}_m\mathbf{A}, \mathbf{C}_1 = \mathbf{Y}_1 = \mathbf{E}_1 = \mathbf{T}_1 = \mathbf{Z}_1 = \Theta_1 = \mathbf{0}_{m \times n} \in \mathbb{R}^{m \times n}, \Psi = \mathbf{0}_m, \operatorname{corr}^{aa} = \operatorname{sum}(\hat{\mathbf{A}}^* \odot \hat{\mathbf{A}}, 2), \lambda, \rho_1, \gamma > 1, q$.

Output: \mathbf{c}

while not converged **do**

compute: $\mathbf{E}_{k+1} = \rho_k\mathbf{Z}_k - \rho_k\Theta_k - \mathbf{C}_k$

$\hat{\Psi}^*$ update: $\hat{\Psi}_{k+1}^* =$

$(\operatorname{sum}(\hat{\mathbf{A}}^* \odot \hat{\mathbf{E}}^*, 2) - \hat{\mathbf{b}}^*) / (\rho_k \operatorname{corr}^{aa} + \frac{1}{2})$

compute: $\tilde{\mathbf{A}}^H\Psi_{k+1}$, see Eq (15)

Θ update: $(\Theta_{k+1}) = (\mathbf{Z} - \frac{1}{\rho_k}\mathbf{C}_k - \tilde{\mathbf{A}}^H\Psi_{k+1})$

$\Theta_{(1:\text{end},:)} = \mathbf{0}_n^H$

compute: $\mathbf{T}_{k+1} = \tilde{\mathbf{A}}^H\Psi_{k+1} + \Theta_{k+1} + \mathbf{C}_k/\rho_k$

\mathbf{Z} update: $\mathbf{Z}_{k+1} = \operatorname{sign}(\mathbf{T}_{k+1}) \odot \min(|\mathbf{T}_{k+1}|, \lambda)$

\mathbf{C} update:

$\mathbf{C}_{k+1} = \mathbf{Y}_k + \rho_k(\tilde{\mathbf{A}}^H\Psi_{k+1} + \Theta_{k+1} - \mathbf{Z}_{k+1})$

compute: $\mathbf{Y}_{k+1} = (1+q)\mathbf{C}_{k+1} - q\mathbf{C}_k$

$\rho_{k+1} = \gamma\rho_k$

end

$\mathbf{c} \leftarrow \mathbf{C}(1:\text{end},:)$

to reshape any of the variables for any operation unlike the case in Algorithm 1.

Regarding the memory efficiency, often there are overhead in computing the FFT of many long vectors that is when both (m, n) are large. For instance, computing $\mathbf{F}\mathbf{E}$ for large (m, n) . The overhead of storage due to replicating \mathbf{E} is not always done efficiently. To circumvent this issue, we simply split the matrix \mathbf{E} into two pieces horizontally $\mathbf{E}_1, \mathbf{E}_2 \in \mathbb{R}^{\frac{m}{2} \times n}$ into an even and odd fashion similar to the butterfly algorithm used to compute the FFTs. Now we only compute the FFT of smaller vectors $\frac{m}{2}$ where we can simply combine them later. For completeness we show here how the procedure is carried out.

Let $\mathbf{v} \in \mathbb{R}^m$, it can be written as follows $\mathbf{v} = \mathbf{v}_1 + \mathbf{v}_2$ where $\mathbf{v}_1 = [v_0 \ 0 \ v_2 \ \dots]^H$ and where $\mathbf{v}_2 = [0 \ v_1 \ 0 \ v_3 \ \dots]^H$. It is obvious that $\tilde{\mathbf{v}} = \tilde{\mathbf{v}}_1 + \tilde{\mathbf{v}}_2$ where $\tilde{\mathbf{v}}$ is the fourier transform of \mathbf{v} . All is left is to com-

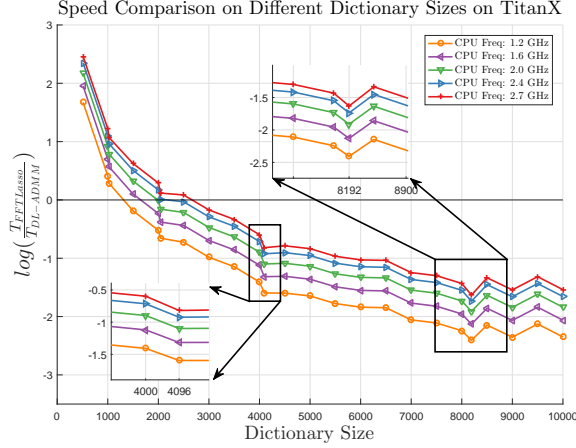

(a) $m = n$.(b) $m = n$.

Figure 1. Runtime comparison between FFTLasso-GPU (on TitanX) and DL-ADMM (multi-core) at different CPU frequencies for square dictionaries ($m = n$)

pute both $\tilde{\mathbf{v}}_1$, and $\tilde{\mathbf{v}}_2$ using lower order FFT as follows:

$$\begin{aligned}
 \mathbf{F}\mathbf{v}_1 &= \sum_{i=0}^m \mathbf{v}_1(i) \exp \frac{-j2\pi ki}{m} = \sum_{0,2,4,\dots}^m \mathbf{v}_1(i) \exp \frac{-j2\pi ki}{m} \\
 &= \sum_{i=0,2,4,\dots}^m \mathbf{v}_e\left(\frac{i}{2}\right) \exp \frac{-j2\pi ki}{m} = \sum_{v=0,1,\dots}^{\frac{m}{2}} \mathbf{v}_e(v) \exp \frac{-j2\pi kv}{\frac{m}{2}} \\
 &= \tilde{\mathbf{v}}_e(i) \forall i = 0, 1, \dots, m
 \end{aligned}$$

That means $\tilde{\mathbf{v}}_1 = [\tilde{\mathbf{v}}_e \ \tilde{\mathbf{v}}_e]^H$ which is only a repeated version of the FFT of $\frac{m}{2}$ of the even part of \mathbf{v} and where $\mathbf{v}_e = [v_0 \ v_2 \ \dots]$. A similar trick can be applied to \mathbf{v}_2 followed by a shift to the right. One can easily show that $\tilde{\mathbf{v}}_1 = [\tilde{\mathbf{v}}_e \ \tilde{\mathbf{v}}_e]^H \odot \mathbf{p}$ where $\mathbf{p}(i) = \exp \frac{-j2\pi ki}{m} \ \forall i = 0, 1, \dots, m$. Therefore,

$$\tilde{\mathbf{v}} = [\tilde{\mathbf{v}}_e^H \ \tilde{\mathbf{v}}_e^H]^H + [\tilde{\mathbf{v}}_o^H \ \tilde{\mathbf{v}}_o^H]^H \odot \mathbf{p}$$

Lastly, in here we present the results of both the TitanX (1.0GHz/core) and TitanX Pascal (1.4GHz/core) experiments comparing FFTLasso to different multi-core CPU frequencies of DL-ADMM. Figure 1, shows a comparison between several hardware (GPUs) of FFTLasso against DL-ADMM on high frequency multi core CPUs. It is clear that depending on the hardware used and the dictionary \mathbf{A} used one would prefer using DL-ADMM over FFTLasso and vice versa. However, figure 1 clearly shows incremental monotone improvement in the speed of FFTLasso as the dictionary size grows large. Specifically, FFTLasso is mostly optimized for any dictionary that is of powers of an integer radix (in this case radix 2).