

# Appendices

Table 4: Symbolic notation for EigenPro 3.0 in Algorithm 1. They satisfy  $m < n$ , and  $q < s < n$ .

Symbol	Purpose
$n$	Number of samples
$m$	Batch-size
$p$	Model size
$s$	Nyström approximation subsample size
$q$	Preconditioner level

## A PROOFS OF INTERMEDIATE RESULTS

### A.1 PROOF OF PROPOSITION 1

**Proposition** (Nyström extension). For  $1 \leq i \leq n$ , let  $\lambda_i$  be an eigenvalue of  $\mathcal{K}$ , and  $\psi_i$  its unit  $\mathcal{H}$ -norm eigenfunction, i.e.,  $\mathcal{K}\{\psi_i\} = \lambda_i\psi_i$ . Then  $\lambda_i$  is also an eigenvalue of  $K(X, X)$ . Moreover if  $\mathbf{e}_i$  is its unit-norm eigenvector, i.e.,  $K(X, X)\mathbf{e}_i = \lambda_i\mathbf{e}_i$ , we have,

$$\psi_i = K(\cdot, X) \frac{\mathbf{e}_i}{\sqrt{\lambda_i}}. \quad (37)$$

*Proof.* Let  $\psi \in \mathcal{H}$  be an eigenfunction of  $\mathcal{K}$ . Then by definition of  $\mathcal{K}$  we have,

$$\lambda\psi = \mathcal{K}\{\psi\} = \sum_{i=1}^n K(\cdot, \mathbf{x}_i)\psi(\mathbf{x}_i). \quad (38)$$

As the result we can write  $\psi$  as below,

$$\psi = \sum_{i=1}^n \frac{\psi(\mathbf{x}_i)}{\lambda} K(\cdot, \mathbf{x}_i). \quad (39)$$

If we apply covariance operator to the both side of 39 we have,

$$\mathcal{K}\{\psi\} = \mathcal{K}\left\{\sum_{i=1}^n \frac{\psi(\mathbf{x}_i)}{\lambda} K(\cdot, \mathbf{x}_i)\right\} = \sum_{i,j=1}^n \frac{\psi(\mathbf{x}_i)}{\lambda} K(\mathbf{x}_i, \mathbf{x}_j)K(\cdot, \mathbf{x}_j) = \sum_{j=1}^n \psi(\mathbf{x}_j)K(\cdot, \mathbf{x}_j). \quad (40)$$

The last equation hold because of equation (38). If we define vector  $\beta$  such that  $\beta_i = \frac{\psi(\mathbf{x}_i)}{\lambda}$ , then 40 can be rewritten as,

$$\sum_{i=1}^n \sum_{j=1}^n \beta_i K(\mathbf{x}_i, \mathbf{x}_j)K(\cdot, \mathbf{x}_j) = \lambda \sum_{i=1}^n \beta_i K(\cdot, \mathbf{x}_i). \quad (41)$$

Compactly we can write 41 as below,

$$K(X, X)^2\beta = \lambda K(X, X)\beta \implies K(X, X)\beta = \lambda\beta.$$

The last implication holds because  $K(X, X)$  is invertable. Thus  $\beta$  is an eigenvector of  $K(X, X)$ . It remains to determine the scale of  $\beta$ .

Now, norm of  $\psi$  can be simplified as

$$\|\psi\|_{\mathcal{H}}^2 = \left\langle \sum_{i=1}^n \beta_i K(\cdot, \mathbf{x}_i), \sum_{j=1}^n \beta_j K(\cdot, \mathbf{x}_j) \right\rangle_{\mathcal{H}} \quad (42)$$

$$= \sum_{i,j=1}^n \beta_i \beta_j \langle K(\cdot, \mathbf{x}_i), K(\cdot, \mathbf{x}_j) \rangle_{\mathcal{H}} = \beta^\top K(X, X)\beta = \lambda \|\beta\|^2. \quad (43)$$

Since  $\psi$  is unit norm, we have  $\|\beta\| = \frac{1}{\sqrt{\lambda}}$ . This concludes the proof.  $\square$

## A.2 PROOF OF LEMMA 2

**Lemma** (Nyström preconditioning). Let  $\mathbf{a} \in \mathbb{R}^m$ , then we have that,

$$\mathcal{P}_s \{K(\cdot, X_m)\mathbf{a}\} = K(\cdot, X_m)\mathbf{a} - K(\cdot, X_s)\mathbf{Q}_s K(X_s, X_m)\mathbf{a}. \quad (44)$$

Where  $\mathbf{Q}_s = E_{s,q}(\mathbf{I}_n - \lambda_{s,q+1}\Lambda_{s,q}^{-1})\Lambda_{s,q}^{-1}E_{s,q}^\top$ .

*Proof.* Recall that  $\mathcal{P}_s := \mathcal{I} - \sum_{i=1}^q \left(1 - \frac{\lambda_{q+1}}{\lambda_i}\right) \psi_i \otimes \psi_i$ . By this definition we can write,

$$\begin{aligned} \mathcal{P}_s (K(\cdot, X_M)\boldsymbol{\alpha}) &= K(\cdot, X_M)\boldsymbol{\alpha} - \sum_{i=1}^s \left(1 - \frac{\lambda_{q+1}}{\lambda_i^s}\right) \langle \psi_i^s, K(\cdot, X_M)\boldsymbol{\alpha} \rangle_{\mathcal{H}} \psi_i^s \\ &= K(\cdot, X_M)\boldsymbol{\alpha} - \sum_{i=1}^q \frac{1}{\lambda_i^s} \left(1 - \frac{\lambda_{q+1}}{\lambda_i^s}\right) \langle K(\cdot, X_s)\mathbf{e}_i, K(\cdot, X_M)\boldsymbol{\alpha} \rangle_{\mathcal{H}} K(\cdot, X_s)\mathbf{e}_i \\ &= K(\cdot, X_M)\boldsymbol{\alpha} - \sum_{i=1}^q \frac{1}{\lambda_i^s} \left(1 - \frac{\lambda_{q+1}}{\lambda_i}\right) \langle K(\cdot, X_s)\mathbf{e}_i, K(\cdot, X_M)\boldsymbol{\alpha} \rangle_{\mathcal{H}} K(\cdot, X_s)\mathbf{e}_i \\ &= K(\cdot, X_M)\boldsymbol{\alpha} - \sum_{i=1}^q \left(1 - \frac{\lambda_{q+1}}{\lambda_i^s}\right) K(\cdot, X_s)\mathbf{e}_i \mathbf{e}_i^\top K(X_s, X_M)\boldsymbol{\alpha}. \end{aligned}$$

Note that we used proposition 1 for  $\psi$ . Now we can compactly write the last expression as below,

$$\begin{aligned} \mathcal{P}_s (K(\cdot, X_M)\boldsymbol{\alpha}) &= K(\cdot, X_M)\boldsymbol{\alpha} - K(\cdot, X_s)E_{s,q}(\mathbf{I}_n - \lambda_{s,q+1}\Lambda_{s,q}^{-1})\Lambda_{s,q}^{-1}E_{s,q}^\top K(X_s, X_M)\boldsymbol{\alpha} \\ &= K(\cdot, X_M)\boldsymbol{\alpha} - K(\cdot, X_s)\mathbf{Q}_s K(X_s, X_M)\boldsymbol{\alpha}. \end{aligned}$$

This concludes the proof. □

## B DETAILS ON EigenPro 2.0

**Lemma 3.** *The iteration in  $\mathbb{R}^n$*

$$\boldsymbol{\alpha}^{t+1} = \boldsymbol{\alpha}^{t+1} - \eta(\mathbf{I}_n - \mathbf{Q})(K(X, X)\boldsymbol{\alpha}^t - \mathbf{y}), \quad (45)$$

where  $\mathbf{Q} = \mathbf{E}(\mathbf{I}_n - \lambda_{q+1}\Lambda_q^{-1})\mathbf{E}^\top$ , emulates the following iteration in  $\mathcal{H}$ .

$$f^{t+1} = f^t - \eta \mathcal{P} \{ \nabla_f L(f^t) \}. \quad (46)$$

*Proof.* Recall that  $\nabla_f L(f^t) = K(\cdot, X)(f^t(X) - \mathbf{y})$  from equation (10), and  $f^t(X) = K(X, X)\boldsymbol{\alpha}^t$  from equation (19). We define  $\mathbf{g}^t := f^t(X) - \mathbf{y} = K(X, X)\boldsymbol{\alpha}^t - \mathbf{y}$ . Following steps of the proof in Appendix A.2 we have

$$\begin{aligned} \mathcal{P} \{ \nabla_f L(f^t) \} &= K(\cdot, X)\mathbf{g}^t - \sum_{i=1}^q \left(1 - \frac{\lambda_{q+1}}{\lambda_i}\right) K(\cdot, X)\mathbf{e}_i \mathbf{e}_i^\top K(X, X)\mathbf{g}^t \\ &= K(\cdot, X)\mathbf{g}^t - K(\cdot, X)\mathbf{E}(\mathbf{I}_n - \lambda_{q+1}\Lambda_q^{-1})\Lambda^{-1}\mathbf{E}^\top K(X, X)\mathbf{g}^t \\ &\stackrel{(a)}{=} K(\cdot, X)\mathbf{g}^t - K(\cdot, X)\mathbf{E}(\mathbf{I}_n - \lambda_{q+1}\Lambda_q^{-1})\Lambda^{-1}\mathbf{E}^\top \mathbf{E}\Lambda\mathbf{E}^\top \mathbf{g}^t \\ &= K(\cdot, X)\mathbf{g}^t - K(\cdot, X)\mathbf{E}(\mathbf{I}_n - \lambda_{q+1}\Lambda_q^{-1})\mathbf{E}^\top \mathbf{g}^t \\ &= K(\cdot, X)\mathbf{g}^t - K(\cdot, X)\mathbf{Q}\mathbf{g}^t \\ &= K(\cdot, X)(\mathbf{I}_n - \mathbf{Q})\mathbf{g}^t. \end{aligned}$$

---

**Algorithm 3** EigenPro 2.0( $X, \mathbf{y}$ ). Solves the linear system  $K(X, X)\boldsymbol{\theta} = \mathbf{y}$

---

**Require:** Data  $(X, \mathbf{y})$ , Nyström size  $s$ , preconditioner level  $q$   
 $\boldsymbol{\alpha} \leftarrow \mathbf{0} \in \mathbb{R}^n$  ▷ initialization  
 $X_s, (\mathbf{E}, \mathbf{D}), \lambda_{q+1}, m \leftarrow \text{EigenPro 2.0\_setup}(X, s, q)$   
Set batchsize  $m \leftarrow \frac{1}{\lambda_{q+1}}$   
**while** Stopping criterion not reached **do**  
 $\boldsymbol{\alpha} \leftarrow \text{EigenPro 2.0\_iteration}(X, \mathbf{y}, X_s, \mathbf{E}, \mathbf{D}, \boldsymbol{\alpha}, m, \eta)$   
**end while**  
**return**  $\boldsymbol{\alpha}$

EigenPro2\_setup( $X, s, q$ )

**Require:** Data  $X$ , Nyström size  $s$ , preconditioner size  $q$   
Fetch a subsample  $X_s \subseteq X$  of size  $s$   
 $(\mathbf{E}, \Lambda) \leftarrow \text{top-}q \text{ eigensystem of } K(X_s, X_s)$  ▷  $\mathbf{E} \in \mathbb{R}^{q \times s}, \Lambda = \text{diag}(\lambda_i) \in \mathbb{R}^{q \times q}$   
 $\mathbf{D}_{ii} = \frac{1}{s\lambda_i} \left(1 - \frac{\lambda_{q+1}}{\lambda_i}\right)$   
 $\beta \leftarrow \max_i K(\mathbf{x}_i, \mathbf{x}_i) \in S$   
 $m \leftarrow \min\left(\frac{\beta}{\lambda_{q+1}}, \text{bs}_{\text{gpu}}\right)$  ▷ batch size<sup>2</sup>  
 $\eta \leftarrow \begin{cases} \frac{\beta}{2m} & m < \frac{\beta}{\lambda_{q+1}} \\ \frac{0.99m}{\beta + (m-1)\lambda_{q+1}} & \text{otherwise} \end{cases}$  ▷ learning rate  
**return**  $X_s, (\mathbf{E}, \mathbf{D}), \eta, m$

EigenPro2\_iteration( $X, \mathbf{y}, X_s, \mathbf{E}, \mathbf{D}, \boldsymbol{\alpha}, m, \eta$ )

**Require:** Data  $(X, \mathbf{y})$ , Nyström subset  $X_s$ , preconditioner  $(\mathbf{E}, \mathbf{D})$ , current estimate  $\boldsymbol{\alpha}$ , batchsize  $m$   
Fetch minibatch  $(X_m, \mathbf{y}_m)$  of size  $m$   
 $\mathbf{g}_m \leftarrow K(X_m, X)\boldsymbol{\alpha} - \mathbf{y}_m$  ▷ stochastic gradient  
 $\boldsymbol{\alpha}_m \leftarrow \boldsymbol{\alpha}_m - \frac{\eta}{m}\mathbf{g}_m$  ▷ gradient step  
 $\boldsymbol{\alpha}_s \leftarrow \boldsymbol{\alpha}_s + \mathbf{E}\mathbf{D}\mathbf{E}^\top K(X_s, X_m)\mathbf{g}_m$  ▷ gradient correction  
**return** Updated estimate  $\boldsymbol{\alpha}$

---

Where (a) follows from  $K(X, X) = \mathbf{E}\Lambda\mathbf{E}^\top$ . Now since  $f^t = K(\cdot, X)\boldsymbol{\alpha}^t$ , equation (46) can be rewritten,

$$\begin{aligned} f^{t+1} &= K(\cdot, X)\boldsymbol{\alpha}^{t+1} - \eta K(\cdot, X)(\mathbf{I}_n - Q)\mathbf{g}^t \\ &= K(\cdot, X)(\boldsymbol{\alpha}^{t+1} - \eta(\mathbf{I}_n - Q)\mathbf{g}^t). \end{aligned}$$

Replacing  $\mathbf{g}^t = K(X, X)\boldsymbol{\alpha}^t - \mathbf{y}$  leads to final update rule below,

$$f^{t+1} = K(\cdot, X)(\boldsymbol{\alpha}^{t+1} - \eta(\mathbf{I}_n - Q)(K(X, X)\boldsymbol{\alpha}^t - \mathbf{y})).$$

This concludes the proof. □

Thus each update constitutes a *stochastic gradient step* which consists updating  $m$  weights corresponding to a minibatch size  $m$ , followed by a *gradient correction* which consists of updating all  $n$  weights.

A higher preconditioner level  $q$  also allows for a higher optimal batch size  $m$  and hence better GPU utilization, see [Ma et al. \(2018\)](#) for details.

With this approximation, the gradient correction simplifies drastically, and only  $s$  weights need to be updated.

---

<sup>2</sup>bs<sub>gpu</sub> is the maximum batch-size that the GPU allows.

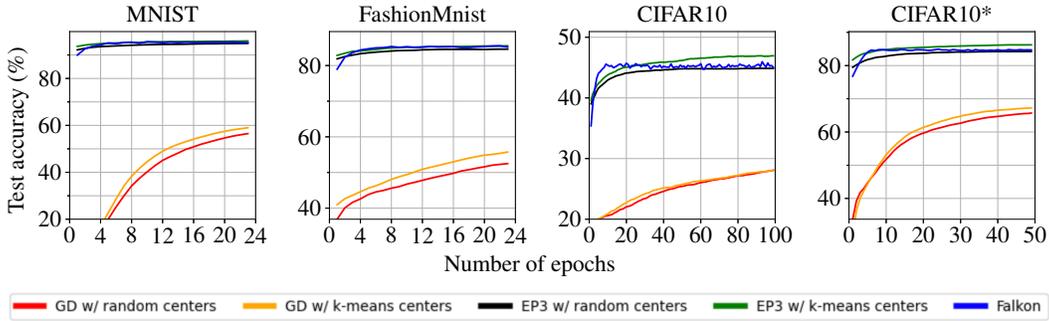


Figure 3: **Comparison with gradient descent and Falkon:** Figure 3 shows the slow convergence of gradient descent given in (50) compared to our algorithm and FALKON from Rudi et al. (2017). Note that FALKON involves a matrix inverse for a projection operation and hence converges faster.

## C DETAILS ON EXPERIMENTS AND IMPLEMENTATION OF ALGORITHM 1

### C.1 COMPUTATIONAL RESOURCES USED

This work used the Extreme Science and Engineering Discovery Environment (XSEDE) (Townes et al., 2014). We used machines with 2x NVIDIA-V100 GPUs, each with a memory of 32GB, and 4x cores of Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz with a RAM of 100 GB.

### C.2 CHOICE OF HYPERPARAMETERS

We choose hyperparameters to minimize computation and maximize GPU utilization. The only hyperparameters that we need to set are  $s, q$  for outer gradient step, and  $\sigma, \xi$  for projection sub-problem. For  $\sigma, \xi$ , we used the same criteria as Ma & Belkin (2019) to optimally use GPU utilization. For  $s, q$ , we prefer larger  $q$  because as it is explained in Ma et al. (2018), larger  $q$  allows for larger learning rate and better condition number. However, in our algorithm we need to approximate the top  $q$  eigensystem of Nyström sub-samples matrix. We used Scipy Virtanen et al. (2020) library to approximate these eigensystem. The stability and precision of these approximations depends on how large is the ratio of  $\frac{s}{q}$ . Empirically we need this ratio to be larger than 10. On the other hand increasing  $s$  will increase setup cost, computation cost and memory cost. We take steps below to choose  $q$  and  $s$ ,

1. We first choose  $s$  as big as our GPU memory allow
2. We choose  $q \approx \frac{s}{10}$
3. We set batch size and learning rate automatically using the *new* top eigenvalue as it is explained in Ma & Belkin (2019) and Ma et al. (2018).

## D CLASSICAL APPROACH TO LEARNING KERNEL NETWORKS WITH GD

If you plug in the form of (4) into (1), we get

$$\underset{\alpha}{\text{minimize}} L(\alpha) = \sum_{i=1}^n L\left(\sum_{j=1}^p K(\mathbf{x}_i, \mathbf{z}_j)\alpha_j, y_i\right) + \lambda \left\langle \sum_{j=1}^p K(\cdot, \mathbf{z}_j), \sum_{j=1}^p K(\cdot, \mathbf{z}_j) \right\rangle_{\mathcal{H}} \quad (47)$$

$$= \sum_{i=1}^n L(\mathbf{I}_n^{(i)} K(X, Z)\alpha, y_i) + \lambda \alpha^\top K(Z, Z)\alpha, \quad (48)$$

where  $\mathbf{I}_n^{(i)}$  is the  $i^{\text{th}}$  row of identity  $\mathbf{I}_n$ . For the square loss this is

$$\underset{\alpha}{\text{minimize}} \|K(X, Z)\alpha - \mathbf{y}\|^2 + \lambda \alpha^\top K(Z, Z)\alpha. \quad (49)$$

Gradient descent on this problem for the square loss yields the update equation,

$$\alpha^{t+1} = \alpha^t - \eta K(Z, X)((K(X, Z)\alpha^t - \mathbf{y}) - \eta \lambda K(Z, Z)\alpha). \quad (50)$$