
On the Spectrum of Random Features Maps of High Dimensional Data

Zhenyu Liao¹ Romain Couillet^{1,2}

Abstract

Random feature maps are ubiquitous in modern statistical machine learning, where they generalize random projections by means of powerful, yet often difficult to analyze nonlinear operators. In this paper, we leverage the “concentration” phenomenon induced by random matrix theory to perform a spectral analysis on the Gram matrix of these random feature maps, here for Gaussian mixture models of simultaneously large dimension and size. Our results are instrumental to a deeper understanding on the interplay of the nonlinearity and the statistics of the data, thereby allowing for a better tuning of random feature-based techniques.

1. Introduction

Finding relevant features is one of the key steps for solving a machine learning problem. To this end, the backpropagation algorithm is probably the best-known method, with which superhuman performances are commonly achieved for specific tasks in applications of computer vision (Krizhevsky et al., 2012) and many others (Schmidhuber, 2015). But data-driven approaches such as the backpropagation method, in addition to being computationally demanding, fail to cope with limited amounts of available training data.

One successful alternative in this regard is the use of “random features”, exploited both in feed-forward neural networks (Huang et al., 2012; Scardapane & Wang, 2017), in large-scale kernel estimation (Rahimi & Recht, 2008; Vedaldi & Zisserman, 2012) and more recently in random sketching schemes (Keriven et al., 2016). Random feature maps consist in projections randomly exploring the set of nonlinear representations of the data, hopefully extracting

features relevant to some given task. The nonlinearities make these representations more mighty but meanwhile theoretically more difficult to analyze and optimize.

Infinitely large random features maps are nonetheless well understood as they result in (asymptotically) equivalent kernels, the most popular example being random Fourier features and their limiting radial basis kernels (Rahimi & Recht, 2008). Beyond those asymptotic results, recent advances in random matrix theory give rise to unexpected simplification on the understanding of the finite-dimensional version of these kernels, i.e., when the data number and size are large but of similar order as the random feature vector size (El Karoui et al., 2010; Couillet et al., 2016). Following the same approach, in this work, we perform a spectral analysis on the Gram matrix of the random feature matrices. This matrix is of key relevance in many associated machine learning methods (e.g., spectral clustering (Ng et al., 2002) and kernel SVM (Schölkopf & Smola, 2002)) and understanding its spectrum casts an indispensable light on their asymptotic performances. In the remainder of the article, we shall constantly consider spectral clustering as a concrete example of application; however, similar analyses can be performed for other types of random feature-based algorithms.

Our contribution is twofold. From a random matrix theory perspective, it is a natural extension of the sample covariance matrix analysis (Silverstein & Bai, 1995) to a nonlinear setting and can also be seen as the generalization of the recent work of (Pennington & Worah, 2017) to a more practical data model. From a machine learning point of view, we describe quantitatively the mutual influence of different nonlinearities and data statistics on the resulting random feature maps. More concretely, based on the ratio of two coefficients from our analysis, commonly used activation functions are divided into three classes: means-oriented, covariance-oriented and balanced, which eventually allows one to choose the activation function with respect to the statistical properties of the data (or task) at hand, with a solid theoretical basis.

We show by experiments that our results, applicable theoretically only to Gaussian mixture data, show an almost perfect match when applied to some real-world datasets. We are thus optimistic that our findings, although restricted to Gaussian assumptions on the data model, can be applied to

¹Laboratoire des Signaux et Systèmes (L2S), CentraleSupélec, Université Paris-Saclay, France; ²G-STATS Data Science Chair, GIPSA-lab, University Grenoble-Alpes, France. Correspondence to: Zhenyu Liao <zhenyu.liao@l2s.centralesupelec.fr>, Romain Couillet <romain.couillet@centralesupelec.fr>.

a larger set of problems beyond strongly structured ones.

Notations: Boldface lowercase (uppercase) characters stand for vectors (matrices), and non-boldface scalars respectively. $\mathbf{1}_T$ is the column vector of ones of size T , and \mathbf{I}_T the $T \times T$ identity matrix. The notation $(\cdot)^\top$ denotes the transpose operator. The norm $\|\cdot\|$ is the Euclidean norm for vectors and the operator norm for matrices.

In the remainder of this article, we introduce the objects of interest and necessary preliminaries in Section 2. Our main results on the spectrum of random feature maps will be presented in Section 3, followed by experiments on two types of classification tasks in Section 4. The article closes on concluding remarks and envisioned extensions in Section 5.

2. Problem Statement and Preliminaries

Let $\mathbf{x}_1, \dots, \mathbf{x}_T \in \mathbb{R}^p$ be independent data vectors, each belonging to one of K distribution classes $\mathcal{C}_1, \dots, \mathcal{C}_K$. Class \mathcal{C}_a has cardinality T_a , for all $a \in \{1, \dots, K\}$. We assume that the data vector \mathbf{x}_i follows a Gaussian mixture model¹, i.e.,

$$\mathbf{x}_i = \boldsymbol{\mu}_a / \sqrt{p} + \boldsymbol{\omega}_i$$

with $\boldsymbol{\omega}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_a/p)$ for some mean $\boldsymbol{\mu}_a \in \mathbb{R}^p$ and covariance $\mathbf{C}_a \in \mathbb{R}^{p \times p}$ of associated class \mathcal{C}_a .

We denote the data matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{p \times T}$ of size T by cascading all \mathbf{x}_i as column vectors. To extract random features, \mathbf{X} is premultiplied by some random matrix $\mathbf{W} \in \mathbb{R}^{n \times p}$ with i.i.d. entries and then applied entry-wise some nonlinear *activation function* $\sigma(\cdot)$ to obtain the random feature matrix $\boldsymbol{\Sigma} \equiv \sigma(\mathbf{W}\mathbf{X}) \in \mathbb{R}^{n \times T}$, whose columns are simply $\sigma(\mathbf{W}\mathbf{x}_i)$ the associated random feature of \mathbf{x}_i .

In this article, we focus on the Gram matrix $\mathbf{G} \equiv \frac{1}{n} \boldsymbol{\Sigma}^\top \boldsymbol{\Sigma}$ of the random features, the entry (i, j) of which is given by

$$\mathbf{G}_{ij} = \frac{1}{n} \sigma(\mathbf{W}\mathbf{x}_i)^\top \sigma(\mathbf{W}\mathbf{x}_j) = \frac{1}{n} \sum_{k=1}^n \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j)$$

with \mathbf{w}_k^\top the k -th row of \mathbf{W} . Note that all \mathbf{w}_k follow the same distribution, so that taking expectation over $\mathbf{w} \equiv \mathbf{w}_k$ of the above equation one results in the average kernel matrix Φ , with the (i, j) entry of which given by

$$\Phi(\mathbf{x}_i, \mathbf{x}_j) = \mathbb{E}_{\mathbf{w}} \mathbf{G}_{ij} = \mathbb{E}_{\mathbf{w}} \sigma(\mathbf{w}^\top \mathbf{x}_i) \sigma(\mathbf{w}^\top \mathbf{x}_j). \quad (1)$$

When the entries of \mathbf{W} follow a standard Gaussian distribution, one can compute the generic form $\Phi(\mathbf{a}, \mathbf{b}) = \mathbb{E}_{\mathbf{w}} \sigma(\mathbf{w}^\top \mathbf{a}) \sigma(\mathbf{w}^\top \mathbf{b})$ by applying the integral trick from (Williams, 1997), for a large set of nonlinear functions $\sigma(\cdot)$ and arbitrary vector \mathbf{a}, \mathbf{b} of appropriate dimension. We list the results for commonly used functions in Table 1.

¹We normalize the data by $p^{-1/2}$ to guarantee that $\|\mathbf{x}_i\| = O(1)$ with high probability when $\|\mathbf{C}_a\| = O(1)$.

Since the Gram matrix \mathbf{G} describes the correlation of data in the *feature space*, it is natural to recenter \mathbf{G} , and thus Φ by pre- and post-multiplying a projection matrix $\mathbf{P} \equiv \mathbf{I}_T - \frac{1}{T} \mathbf{1}_T \mathbf{1}_T^\top$. In the case of Φ , we get

$$\Phi_c \equiv \mathbf{P} \Phi \mathbf{P}.$$

In the recent line of works (Louart et al., 2018; Pennington & Worah, 2017), it has been shown that the large dimensional (large n, p, T) characterization of \mathbf{G} , in particular its eigenspectrum, is fully determined by Φ and the ratio n/p . For instance, by defining the *empirical spectral distribution* of $\mathbf{G}_c = \mathbf{P}\mathbf{G}\mathbf{P}$ as $\rho_{\mathbf{G}_c}(x) \equiv \frac{1}{T} \sum_{i=1}^T \mathbf{1}_{\lambda_i \leq x}(x)$, with $\lambda_1, \dots, \lambda_T$ the eigenvalues of \mathbf{G}_c , it has been shown in (Louart et al., 2018) that, as $n, p, T \rightarrow \infty$, $\rho_{\mathbf{G}_c}(x)$ almost surely converges to a non-random distribution $\rho(x)$, referred to as the *limiting spectral distribution* of \mathbf{G}_c such that

$$\rho(x) = \frac{1}{\pi} \lim_{y \rightarrow 0^+} \int_{-\infty}^x \Im [m(t + iy)] dt.$$

with $m(z)$ the associated Stieltjes transform given by

$$m(z) = \frac{1}{n} \text{tr} \mathbf{Q}(z), \quad \mathbf{Q}(z) \equiv \left(\frac{\Phi_c}{1 + \delta(z)} - z \mathbf{I}_T \right)^{-1}$$

with $\delta(z)$ the unique solution of $\delta(z) = \frac{1}{n} \text{tr}(\Phi_c \mathbf{Q}(z))$.

As a consequence, in the objective of understanding the asymptotic behavior of \mathbf{G}_c as n, p, T are simultaneously large, we shall focus our analysis on Φ_c . To this end, the following assumptions will be needed throughout the paper.

Assumption 1 (Growth rate). As $T \rightarrow \infty$,

- 1) $p/T \rightarrow c_0 \in (0, \infty)$,
- 2) for each $a \in \{1, \dots, K\}$, $T_a/T \rightarrow c_a \in (0, 1)$,
- 3) $\|\boldsymbol{\mu}_a\| = O(1)$,
- 4) let $\mathbf{C}^\circ \equiv \sum_{a=1}^K \frac{T_a}{T} \mathbf{C}_a$ and for $a \in \{1, \dots, K\}$, $\mathbf{C}_a^\circ \equiv \mathbf{C}_a - \mathbf{C}^\circ$, then $\|\mathbf{C}_a\| = O(1)$ and $\text{tr}(\mathbf{C}_a^\circ)/\sqrt{p} = O(1)$,
- 5) for technical convenience we assume in addition that $\tau \equiv \text{tr}(\mathbf{C}^\circ)/p$ converges in $(0, \infty)$.

Assumption 1 ensures that the information of means or covariances is neither too simple nor impossible to be extracted from the data, as investigated in (Couillet et al., 2016).

Let us now introduce the key steps of our analysis. Under Assumption 1, note that for $\mathbf{x}_i \in \mathcal{C}_a$ and $\mathbf{x}_j \in \mathcal{C}_b$, $i \neq j$,

$$\mathbf{x}_i^\top \mathbf{x}_j = \underbrace{\boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j}_{O(p^{-1/2})} + \underbrace{\boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b/p + \boldsymbol{\mu}_a^\top \boldsymbol{\omega}_j/\sqrt{p} + \boldsymbol{\mu}_b^\top \boldsymbol{\omega}_i/\sqrt{p}}_{O(p^{-1})}$$

Table 1. $\Phi(\mathbf{a}, \mathbf{b})$ for different $\sigma(\cdot)$, $\angle(\mathbf{a}, \mathbf{b}) \equiv \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$.

$\sigma(t)$	$\Phi(\mathbf{a}, \mathbf{b})$
t	$\mathbf{a}^\top \mathbf{b}$
$\max(t, 0) \equiv \text{ReLU}(t)$	$\frac{1}{2\pi} \ \mathbf{a}\ \ \mathbf{b}\ \left(\angle(\mathbf{a}, \mathbf{b}) \arccos(-\angle(\mathbf{a}, \mathbf{b})) + \sqrt{1 - \angle(\mathbf{a}, \mathbf{b})^2} \right)$
$ t $	$\frac{2}{\pi} \ \mathbf{a}\ \ \mathbf{b}\ \left(\angle(\mathbf{a}, \mathbf{b}) \arcsin(\angle(\mathbf{a}, \mathbf{b})) + \sqrt{1 - \angle(\mathbf{a}, \mathbf{b})^2} \right)$
$\varsigma_+ \max(t, 0) + \varsigma_- \max(-t, 0)$	$\frac{1}{2}(\varsigma_+^2 + \varsigma_-^2) \mathbf{a}^\top \mathbf{b} + \frac{\ \mathbf{a}\ \ \mathbf{b}\ }{2\pi} (\varsigma_+ + \varsigma_-)^2 \left(\sqrt{1 - \angle(\mathbf{a}, \mathbf{b})^2} - \angle(\mathbf{a}, \mathbf{b}) \arccos(\angle(\mathbf{a}, \mathbf{b})) \right)$
$1_{t>0}$	$\frac{1}{2} - \frac{1}{2\pi} \arccos(\angle(\mathbf{a}, \mathbf{b}))$
$\text{sign}(t)$	$\frac{2}{\pi} \arcsin(\angle(\mathbf{a}, \mathbf{b}))$
$\varsigma_2 t^2 + \varsigma_1 t + \varsigma_0$	$\varsigma_2 \left(2(\mathbf{a}^\top \mathbf{b})^2 + \ \mathbf{a}\ ^2 \ \mathbf{b}\ ^2 \right) + \varsigma_1 \mathbf{a}^\top \mathbf{b} + \varsigma_2 \varsigma_0 (\ \mathbf{a}\ ^2 + \ \mathbf{b}\ ^2) + \varsigma_0^2$
$\cos(t)$	$\exp\left(-\frac{1}{2}(\ \mathbf{a}\ ^2 + \ \mathbf{b}\ ^2)\right) \cosh(\mathbf{a}^\top \mathbf{b})$
$\sin(t)$	$\exp\left(-\frac{1}{2}(\ \mathbf{a}\ ^2 + \ \mathbf{b}\ ^2)\right) \sinh(\mathbf{a}^\top \mathbf{b})$
$\text{erf}(t)$	$\frac{2}{\pi} \arcsin\left(\frac{2\mathbf{a}^\top \mathbf{b}}{\sqrt{(1+2\ \mathbf{a}\ ^2)(1+2\ \mathbf{b}\ ^2)}}\right)$
$\exp(-\frac{t^2}{2})$	$\frac{1}{\sqrt{(1+\ \mathbf{a}\ ^2)(1+\ \mathbf{b}\ ^2) - (\mathbf{a}^\top \mathbf{b})^2}}$

which allows one to perform a Taylor expansion around 0 as $p, T \rightarrow \infty$, to give a reasonable approximation of nonlinear functions of $\mathbf{x}_i^\top \mathbf{x}_j$, such as those appearing in Φ_{ij} (see again Table 1). For $i = j$, one has instead

$$\|\mathbf{x}_i\|^2 = \underbrace{\|\boldsymbol{\omega}_i\|^2}_{O(1)} + \underbrace{\|\boldsymbol{\mu}_a\|^2/p + 2\boldsymbol{\mu}_a^\top \boldsymbol{\omega}_i/\sqrt{p}}_{O(p^{-1})}.$$

From $\mathbb{E}_{\boldsymbol{\omega}_i}[\|\boldsymbol{\omega}_i\|^2] = \text{tr}(\mathbf{C}_a)/p$ it is convenient to further write $\|\boldsymbol{\omega}_i\|^2 = \text{tr}(\mathbf{C}_a)/p + (\|\boldsymbol{\omega}_i\|^2 - \text{tr}(\mathbf{C}_a)/p)$, where $\text{tr}(\mathbf{C}_a)/p = O(1)$ and $\|\boldsymbol{\omega}_i\|^2 - \text{tr}(\mathbf{C}_a)/p = O(p^{-1/2})$. By definition $\tau \equiv \text{tr}(\mathbf{C}^\circ)/p = O(1)$ and exploiting again Assumption 1 one results in,

$$\|\mathbf{x}_i\|^2 = \underbrace{\tau}_{O(1)} + \underbrace{\text{tr}(\mathbf{C}_a^\circ)/p + \|\boldsymbol{\omega}_i\|^2 - \text{tr}(\mathbf{C}_a)/p}_{O(p^{-1/2})} + \underbrace{\|\boldsymbol{\mu}_a\|^2/p + 2\boldsymbol{\mu}_a^\top \boldsymbol{\omega}_i/\sqrt{p}}_{O(p^{-1})}$$

which allows for a Taylor expansion of nonlinear functions of $\|\mathbf{x}_i\|^2$ around τ , as has been done for $\mathbf{x}_i^\top \mathbf{x}_j$.

From Table 1, it appears that, for every listed $\sigma(\cdot)$, $\Phi(\mathbf{x}_i, \mathbf{x}_j)$ is a smooth function of $\mathbf{x}_i^\top \mathbf{x}_j$ and $\|\mathbf{x}_i\|$, $\|\mathbf{x}_j\|$, despite their possible discontinuities (e.g., the ReLU function and $\sigma(t) = |t|$). The above results thus allow for an entry-wise Taylor expansion of the matrix Φ in the large p, T limit.

A critical aspect of the analysis where random matrix theory comes into play now consists in developing Φ as a sum of matrices arising from the Taylor expansion and ignoring terms that give rise to a vanishing operator norm,

so as to find an asymptotic equivalent matrix $\tilde{\Phi}$ such that $\|\Phi - \tilde{\Phi}\| \rightarrow 0$ as $p, T \rightarrow \infty$, as described in detail in the following section. This analysis provides a simplified asymptotically equivalent expression for Φ with all nonlinearities removed, which is the crux of the present study.

3. Main Results

In the remainder of this article, we shall use the following notations for random elements,

$$\boldsymbol{\Omega} \equiv [\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_T], \boldsymbol{\phi} \equiv \{\|\boldsymbol{\omega}_i\|^2 - \mathbb{E}\|\boldsymbol{\omega}_i\|^2\}_{i=1}^T$$

such that $\boldsymbol{\Omega} \in \mathbb{R}^{p \times T}$, $\boldsymbol{\phi} \in \mathbb{R}^T$. For deterministic elements²,

$$\mathbf{M} \equiv [\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K] \in \mathbb{R}^{p \times K}, \mathbf{t} \equiv \{\text{tr} \mathbf{C}_a^\circ / \sqrt{p}\}_{a=1}^K$$

$$\mathbf{J} \equiv [\mathbf{j}_1, \dots, \mathbf{j}_K] \in \mathbb{R}^{T \times K}, \mathbf{S} \equiv \{\text{tr}(\mathbf{C}_a \mathbf{C}_b) / p\}_{a,b=1}^K$$

with $\mathbf{t} \in \mathbb{R}^K$, $\mathbf{S} \in \mathbb{R}^{K \times K}$ and $\mathbf{j}_a \in \mathbb{R}^T$ denotes the canonical vector of class \mathcal{C}_a such that $(\mathbf{j}_a)_i = \delta_{\mathbf{x}_i \in \mathcal{C}_a}$.

Theorem 1 (Asymptotic Equivalent of Φ_c). *Let Assumption 1 hold and Φ_c be defined as $\Phi_c \equiv \mathbf{P} \Phi \mathbf{P}$, with Φ given in (1). Then, as $T \rightarrow \infty$, for all $\sigma(\cdot)$ given in Table 1,³*

$$\|\Phi_c - \tilde{\Phi}_c\| \rightarrow 0$$

²As a reminder here, \mathbf{M} stands for *means*, \mathbf{t} accounts for (difference in) *traces* while \mathbf{S} for the “*shapes*” of covariances.

³For all functions $\sigma(\cdot)$ listed in Table 1 we identified a “pattern” in the structure of $\tilde{\Phi}_c$, which then led to Theorem 1 and Table 2. This two-step approach does not yet allow to justify whether this pattern goes beyond these (listed) functions; hence Theorem 1 is stated so far solely for these functions.

almost surely, with $\tilde{\Phi}_c = \mathbf{P}\tilde{\Phi}\mathbf{P}$ and

$$\tilde{\Phi} \equiv d_1 \left(\boldsymbol{\Omega} + \mathbf{M} \frac{\mathbf{J}^\top}{\sqrt{p}} \right)^\top \left(\boldsymbol{\Omega} + \mathbf{M} \frac{\mathbf{J}^\top}{\sqrt{p}} \right) + d_2 \mathbf{U} \mathbf{B} \mathbf{U}^\top + d_0 \mathbf{I}_T$$

where we recall that $\mathbf{P} \equiv \mathbf{I}_T - \frac{1}{T} \mathbf{1}_T \mathbf{1}_T^\top$ and

$$\mathbf{U} \equiv \left[\frac{\mathbf{J}}{\sqrt{p}}, \boldsymbol{\phi} \right], \quad \mathbf{B} \equiv \begin{bmatrix} \mathbf{t} \mathbf{t}^\top + 2\mathbf{S} & \mathbf{t} \\ \mathbf{t}^\top & 1 \end{bmatrix}$$

with the coefficients d_0, d_1, d_2 given in Table 2.

We refer the readers to Section A in Supplementary Material for a detailed proof of Theorem 1.

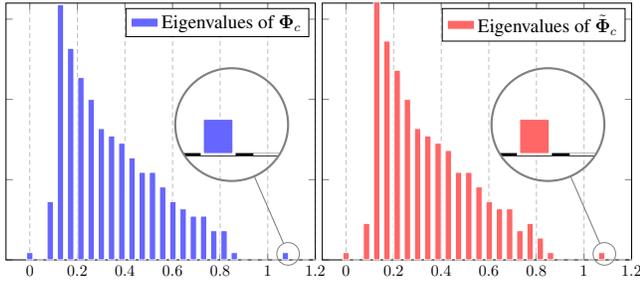


Figure 1. Eigenvalue distribution of Φ_c and $\tilde{\Phi}_c$ for the ReLU function and Gaussian mixture data with $\boldsymbol{\mu}_a = [\mathbf{0}_{a-1}; 3; \mathbf{0}_{p-a}]$, $\mathbf{C}_a = (1 + 2(a-1)/\sqrt{p}) \mathbf{I}_p$, $p = 512$, $T = 256$ and $c_1 = c_2 = 1/2$. Expectation estimated with 500 realizations of \mathbf{W} .

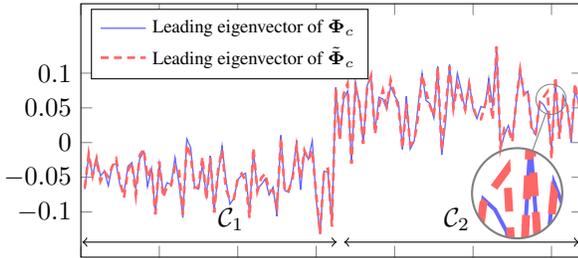


Figure 2. Leading eigenvector of Φ_c and $\tilde{\Phi}_c$ in the settings of Figure 1, with $\mathbf{j}_1 = [\mathbf{1}_{T_1}; \mathbf{0}_{T_2}]$ and $\mathbf{j}_2 = [\mathbf{0}_{T_1}; \mathbf{1}_{T_2}]$.

Theorem 1 tells us as a corollary (see Corollary 4.3.15 in (Horn & Johnson, 2012)) that the maximal difference between the eigenvalues of Φ_c and $\tilde{\Phi}_c$ vanishes asymptotically as $p, T \rightarrow \infty$, as confirmed in Figure 1. Similarly the distance between the “isolated eigenvectors⁴” also vanishes, as seen in Figure 2. This is of tremendous importance as the determination of the leading eigenvalues and eigenvectors of $\tilde{\Phi}_c$ (that contain crucial information for clustering, for example) can be studied from the equivalent problem performed on Φ_c and becomes mathematically more tractable.

⁴Eigenvectors that correspond to the eigenvalues found at a non-vanishing distance from the other eigenvalues.

On closer inspection of Theorem 1, the matrix $\tilde{\Phi}$ is expressed as the sum of three terms, weighted respectively by the three coefficients d_0, d_1 and d_2 , that depend on the nonlinear function $\sigma(\cdot)$ via Table 2. Note that the statistical structure of the data $\{\mathbf{x}_i\}_{i=1}^T$ (namely the means in \mathbf{M} and the covariances in \mathbf{t} and \mathbf{S}) is perturbed by random fluctuations ($\boldsymbol{\Omega}$ and $\boldsymbol{\phi}$) and it is thus impossible to get rid of these noisy terms by wisely choosing the function $\sigma(\cdot)$. This is in sharp contrast to (Couillet et al., 2016) where it is shown that more general kernels (i.e., not arising from random feature maps) allow for a more flexible treatment of information versus noise.

However, there does exist a balance between the means and covariances, that provides some instructions in the appropriate choice of the nonlinearity. From Table 2, the functions $\sigma(\cdot)$ can be divided into the following three groups:

- *mean-oriented*, where $d_1 \neq 0$ while $d_2 = 0$: this is the case of the functions t , $1_{t>0}$, $\text{sign}(t)$, $\sin(t)$ and $\text{erf}(t)$, which asymptotically track only the difference in means (i.e., \mathbf{t} and \mathbf{S} disappear from the expression of $\tilde{\Phi}_c$). As an illustration, in Figure 3 one fails to separate two Gaussian datasets of common mean but of different covariances with the erf function, while ReLU is able to accomplish the task;
- *covariance-oriented*, where $d_1 = 0$ while $d_2 \neq 0$: this concerns the functions $|t|$, $\cos(t)$ and $\exp(-t^2/2)$, which asymptotically track only the difference in covariances. Figure 4 illustrates the impossibility to classify Gaussian mixture with same covariance with $\sigma(t) = |t|$, in contrast to the ReLU function;
- *balanced*, where both $d_1, d_2 \neq 0$: here for the ReLU function $\max(t, 0)$, the Leaky ReLU function (Maas et al., 2013) $\varsigma_+ \max(t, 0) + \varsigma_- \max(-t, 0)$ and the quadratic function $\varsigma_2 t^2 + \varsigma_1 t + \varsigma_0$.

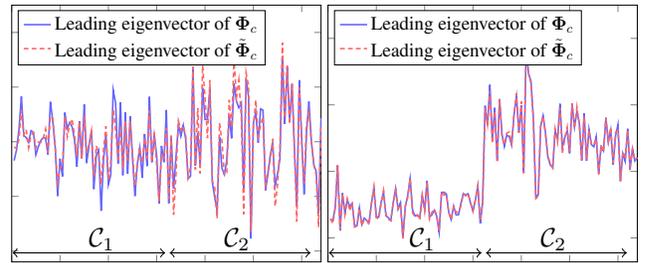


Figure 3. Leading eigenvector of Φ_c and $\tilde{\Phi}_c$ for erf (left) and the ReLU (right) function, performed on Gaussian mixture data with $\boldsymbol{\mu}_a = \mathbf{0}_p$, $\mathbf{C}_a = (1 + 15(a-1)/\sqrt{p}) \mathbf{I}_p$, $p = 512$, $T = 256$, $c_1 = c_2 = 1/2$; $\mathbf{j}_1 = [\mathbf{1}_{T_1}; \mathbf{0}_{T_2}]$ and $\mathbf{j}_2 = [\mathbf{0}_{T_1}; \mathbf{1}_{T_2}]$. Expectation estimated by averaging over 500 realizations of \mathbf{W} .

Table 2. Coefficients d_i in $\tilde{\Phi}_c$ for different $\sigma(\cdot)$.

$\sigma(t)$	d_0	d_1	d_2
t	0	1	0
$\max(t, 0) \equiv \text{ReLU}(t)$	$(\frac{1}{4} - \frac{1}{2\pi})\tau$	$\frac{1}{4}$	$\frac{1}{8\pi\tau}$
$ t $	$(1 - \frac{2}{\pi})\tau$	0	$\frac{1}{2\pi\tau}$
$\varsigma_+ \max(t, 0) + \varsigma_- \max(-t, 0) \equiv \text{LReLU}(t)$	$\frac{\pi-2}{4\pi}(\varsigma_+ + \varsigma_-)^2\tau$	$\frac{1}{4}(\varsigma_+ - \varsigma_-)^2$	$\frac{1}{8\pi\tau}(\varsigma_+ + \varsigma_-)^2$
$1_{t>0}$	$\frac{1}{4} - \frac{1}{2\pi}$	$\frac{1}{2\pi\tau}$	0
$\text{sign}(t)$	$1 - \frac{2}{\pi}$	$\frac{2}{\pi\tau}$	0
$\varsigma_2 t^2 + \varsigma_1 t + \varsigma_0$	$2\tau^2 \varsigma_2^2$	ς_1^2	ς_2^2
$\cos(t)$	$\frac{1}{2} + \frac{e^{-2\tau}}{2} - e^{-\tau}$	0	$\frac{e^{-\tau}}{4}$
$\sin(t)$	$\frac{1}{2} - \frac{e^{-2\tau}}{2} - \tau e^{-\tau}$	$e^{-\tau}$	0
$\text{erf}(t)$	$\frac{2}{\pi} \left(\arccos\left(\frac{2\tau}{2\tau+1}\right) - \frac{2\tau}{2\tau+1} \right)$	$\frac{4}{\pi} \frac{1}{2\tau+1}$	0
$\exp(-\frac{t^2}{2})$	$\frac{1}{\sqrt{2\tau+1}} - \frac{1}{\tau+1}$	0	$\frac{1}{4(\tau+1)^3}$

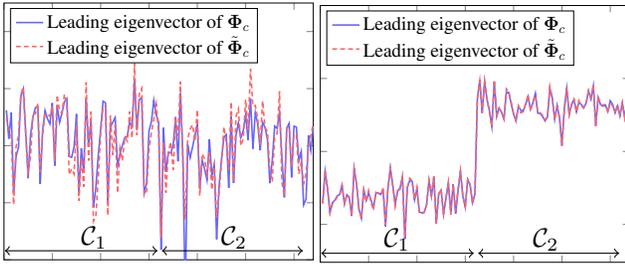


Figure 4. Leading eigenvector of Φ_c and $\tilde{\Phi}_c$ for $\sigma(t) = |t|$ (left) and the ReLU (right) function, performed on Gaussian mixture data with $\mu_a = [\mathbf{0}_{a-1}; 5; \mathbf{0}_{p-a}]$, $\mathbf{C}_a = \mathbf{I}_p$, $p = 512$, $T = 256$, $c_1 = c_2 = 1/2$, $\mathbf{j}_1 = [\mathbf{1}_{T_1}; \mathbf{0}_{T_2}]$ and $\mathbf{j}_2 = [\mathbf{0}_{T_1}; \mathbf{1}_{T_2}]$. Expectation estimated by averaging over 500 realizations of \mathbf{W} .

Before entering into a more detailed discussion of Theorem 1, first note importantly that, for practical interests, the quantity τ can be estimated consistently from the data, as described in the following lemma.

Lemma 1 (Consistent estimator of τ). *Let Assumption 1 hold and recall the definition $\tau \equiv \text{tr}(\mathbf{C}^\circ) / p$. Then, as $T \rightarrow \infty$, with probability 1*

$$\frac{1}{T} \sum_{i=1}^T \|\mathbf{x}_i\|^2 - \tau \rightarrow 0.$$

Proof. Since

$$\frac{1}{T} \sum_{i=1}^T \|\mathbf{x}_i\|^2 = \frac{1}{T} \sum_{a=1}^K \sum_{i=1}^{T_a} \frac{1}{p} \|\mu_a\|^2 - \frac{2}{\sqrt{p}} \mu_a^\top \omega_i + \|\omega_i\|^2,$$

with Assumption 1 we have $\frac{1}{T} \sum_{a=1}^K \sum_{i=1}^{T_a} \frac{1}{p} \|\mu_a\|^2 = O(p^{-1})$. The term $\frac{1}{T} \sum_{a=1}^K \sum_{i=1}^{T_a} \frac{2}{\sqrt{p}} \mu_a^\top \omega_i$ is a linear com-

bination of independent zero-mean Gaussian variables and vanishes with probability 1 as $p, T \rightarrow \infty$ with Chebyshev's inequality and the Borel-Cantelli lemma. Ultimately by the strong law of large numbers, we have $\frac{1}{T} \sum_{i=1}^T \|\omega_i\|^2 - \tau \rightarrow 0$ almost surely, which concludes the proof. \square

From a practical aspect, a few remarks on the conclusions of Theorem 1 can be made.

Remark 1 (Constant shift in feature space). *For $\sigma(t) = \varsigma_2 t^2 + \varsigma_1 t + \varsigma_0$, note the absence of ς_0 in Table 2, meaning that the constant of the quadratic function does not affect the spectrum of the feature matrix. More generally, it can be shown through the integral trick of (Williams, 1997) that the function $\sigma(t) + c$ for some constant shift c gives the same matrix Φ_c as the original function $\sigma(t)$.*

A direct consequence of Remark 1 is that the coefficients d_0, d_1, d_2 of the function $\text{sign}(t)$ are four times those of $1_{t>0}$, as a result of the fact that $\text{sign}(t) = 2 \cdot 1_{t>0} - 1$. Constant shifts have, as such, no consequence in classification or clustering applications.

Remark 2 (Universality of quadratic and Leaky ReLU functions). *Ignoring the coefficient d_0 that gives rise to a constant shift of all eigenvalues of $\tilde{\Phi}_c$ and thus of no practical relevance, observe from Table 2 that by tuning the parameters of the quadratic and Leaky ReLU functions (LReLU(t)), one can select arbitrary positive value for the ratio d_1/d_2 , while the other listed functions have constraints linking d_1 to d_2 .*

Following the discussions in Remark 2, the parameters ς_+, ς_- of the LReLU, as well as ς_1, ς_2 of the quadratic function, essentially act to balance the weights of means

and covariances in the mixture model of the data. More precisely, as $\frac{\varsigma_+}{\varsigma_-} \rightarrow 1$ or $\varsigma_2 \gg \varsigma_1$, more emphasis is set on the “distance” between covariance matrices while $\frac{\varsigma_+}{\varsigma_-} \rightarrow -1$ or $\varsigma_1 \gg \varsigma_2$ stresses the differences in means.

In Figure 5, spectral clustering on four classes of Gaussian data is performed: $\mathcal{N}(\boldsymbol{\mu}_1, \mathbf{C}_1)$, $\mathcal{N}(\boldsymbol{\mu}_1, \mathbf{C}_2)$, $\mathcal{N}(\boldsymbol{\mu}_2, \mathbf{C}_1)$ and $\mathcal{N}(\boldsymbol{\mu}_2, \mathbf{C}_2)$ with the LReLU function that takes different values for ς_+ and ς_- . For $a = 1, 2$, $\boldsymbol{\mu}_a = [\mathbf{0}_{a-1}; 5; \mathbf{0}_{p-a}]$ and $\mathbf{C}_a = \left(1 + \frac{15(a-1)}{\sqrt{p}}\right) \mathbf{I}_p$. By choosing $\varsigma_+ = \varsigma_- = 1$ (equivalent to the linear map $\sigma(t) = t$) and $\varsigma_+ = -\varsigma_- = 1$ (equivalent to $\sigma(t) = |t|$), with the leading two eigenvectors we always recover two classes instead of four, as each setting of parameters only allows for a part of the statistical information of the data to be used for clustering. However, by taking $\varsigma_+ = 1, \varsigma_- = 0$ (the ReLU function) we distinguish all four classes in the leading two eigenvectors, to which the k-means method can then be applied for final classification, as shown in Figure 6.

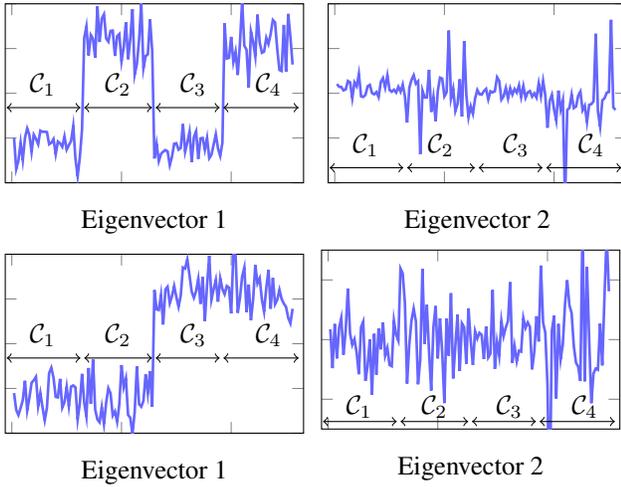


Figure 5. Leading two eigenvectors of Φ_c for the LReLU function with $\varsigma_+ = \varsigma_- = 1$ (top) and $\varsigma_+ = -\varsigma_- = 1$ (bottom), performed on four classes Gaussian mixture data with $p = 512$, $T = 256$, $c_a = \frac{1}{4}$ and $\mathbf{j}_a = [\mathbf{0}_{T_a-1}; \mathbf{1}_{T_a}; \mathbf{0}_{T-T_a}]$, for $a = 1, 2, 3, 4$. Expectation estimated by averaging over 500 realizations of \mathbf{W} .

Of utmost importance for random feature-based spectral methods (such as kernel spectral clustering discussed above (Ng et al., 2002)) is the presence of informative eigenvectors in the spectrum of \mathbf{G} , and thus of Φ_c . To gain a deeper understanding on the spectrum of Φ_c , one can rewrite $\tilde{\Phi}$ in the more compact form,

$$\tilde{\Phi} = d_1 \boldsymbol{\Omega}^\top \boldsymbol{\Omega} + \mathbf{V} \mathbf{A} \mathbf{V}^\top + d_0 \mathbf{I}_T \quad (2)$$

where

$$\mathbf{V} \equiv \left[\frac{\mathbf{J}}{\sqrt{p}}, \boldsymbol{\phi}, \boldsymbol{\Omega}^\top \mathbf{M} \right], \quad \mathbf{A} \equiv \begin{bmatrix} \mathbf{A}_{11} & d_2 \mathbf{t} & d_1 \mathbf{I}_K \\ d_2 \mathbf{t}^\top & d_2 & 0 \\ d_1 \mathbf{I}_K & 0 & 0 \end{bmatrix}$$

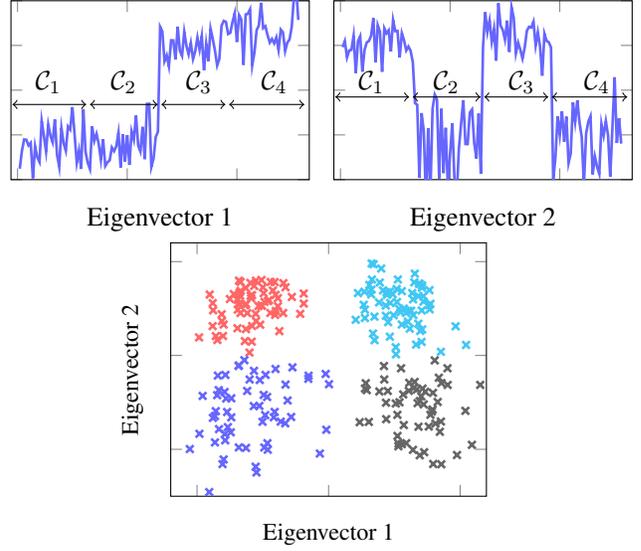


Figure 6. Leading two eigenvectors of Φ_c (top) for the LReLU function with $\varsigma_+ = 1, \varsigma_- = 0$ and two dimensional representation of these eigenvectors (bottom), in the same setting as in Figure 5.

with $\mathbf{A}_{11} \equiv d_1 \mathbf{M}^\top \mathbf{M} + d_2 (\mathbf{t} \mathbf{t}^\top + 2\mathbf{S})$, that is akin to the so-called “spiked model” in the random matrix literature (Baik et al., 2005), as it equals, if $d_1 \neq 0$, the sum of some standard (noise-like) random matrix $\boldsymbol{\Omega}^\top \boldsymbol{\Omega}$, and a low rank (here up to $2K + 1$) informative matrix $\mathbf{V} \mathbf{A} \mathbf{V}^\top$, that may induce some isolated eigenvalues outside the main bulk of eigenvalues in the spectrum of $\tilde{\Phi}_c$, as shown in Figure 1.

The eigenvectors associated to these eigenvalues often contain crucial information about the data statistics (the classes in a classification settings). In particular, note that the matrix \mathbf{V} contains the canonical vector \mathbf{j}_a of class \mathcal{C}_a and we thus hope to find some isolated eigenvector of Φ_c aligned to \mathbf{j}_a that can be directly used to perform clustering. Intuitively speaking, if the matrix \mathbf{A} contains sufficient energy (has sufficiently large operator norm), the eigenvalues associated to the small rank matrix $\mathbf{V} \mathbf{A} \mathbf{V}^\top$ may jump out from the main bulk of $\boldsymbol{\Omega}^\top \boldsymbol{\Omega}$ and becomes “isolated” as in Figure 1, referred to as the *phase transition* phenomenon in the random matrix literature (Baik et al., 2005). The associated eigenvectors then tend to align to linear combinations of the canonical vectors \mathbf{j}_a as seen in Figure 5-6. This alignment between the isolated eigenvectors and \mathbf{j}_a is essentially measured by the amplitude of the eigenvalues of the matrix \mathbf{A}_{11} , or more concretely, the statistical differences of the data (namely, \mathbf{t} , \mathbf{S} and \mathbf{M}). Therefore, a good adaptation of the ratio d_1/d_2 ensures the (asymptotic) detectability of different classes from the spectrum of Φ_c .

4. Numerical Validations

We complete this article by showing that our theoretical results, derived from Gaussian mixture models, show an unexpected close match in practice when applied to some real-world datasets. We consider two different types of classification tasks: one on handwritten digits of the popular MNIST (LeCun et al., 1998) database (number 6 and 8), and the other on epileptic EEG time series data (Andrzejak et al., 2001) (set B and E). These two datasets are typical examples of means-dominant (handwritten digits recognition) and covariances-dominant (EEG times series classification) tasks. This is numerically confirmed in Table 3. Python 3 codes to reproduce the results in this section are available at <https://github.com/Zhenyu-LIAO/RMT4RFM>.

Table 3. Empirical estimation of (normalized) differences in means and covariances of the MNIST (Figure 7 and 8) and epileptic EEG (Figure 9 and 10) datasets.

	$\ \mathbf{M}^T\mathbf{M}\ $	$\ \mathbf{t}\mathbf{t}^T + 2\mathbf{S}\ $
MNIST DATA	172.4	86.0
EEG DATA	1.2	182.7

4.1. Handwritten digits recognition

We perform random feature-based spectral clustering on data matrices that consist of $T = 32, 64$ and 128 randomly selected vectorized images of size $p = 784$ from the MNIST dataset. The “true” means and covariances are empirically obtained from the full set of 11 769 MNIST images (5 918 images of number 6 and 5 851 of number 8) to construct the matrix $\tilde{\Phi}_c$ as per Theorem 1. Comparing the matrix Φ_c built from the data and the theoretically equivalent $\tilde{\Phi}_c$ obtained as if the data were Gaussian with the (empirically) computed means and covariances, we observe an extremely close fit in the behavior of the eigenvalues in Figure 7, as well of the leading eigenvector in Figure 8. The k-means method is then applied to the leading two eigenvectors of the matrix $\mathbf{G}_c \in \mathbb{R}^{T \times T}$ that consists of $n = 32$ random features to perform unsupervised classification, with resulting accuracies (averaged over 50 runs) reported in Table 4. As remarked from Table 3, the mean-oriented $\sigma(t)$ functions are expected to outperform the covariance-oriented ones in this task, which is consistent with the results in Table 4.

4.2. EEG time series classification

The epileptic EEG dataset⁵, developed by the University of Bonn, Germany, is described in (Andrzejak et al., 2001). The dataset consists of five subsets (denoted A-E), each containing 100 single-channel EEG segments of 23.6-sec

⁵<http://www.meb.unibonn.de/epileptologie/science/physik/eegdata.html>.

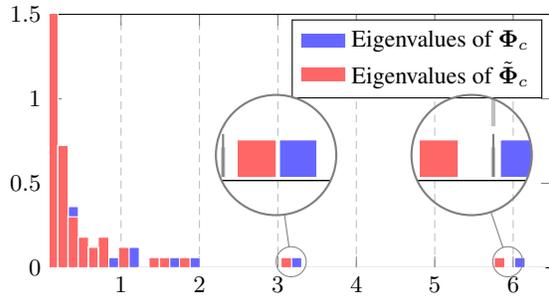


Figure 7. Eigenvalue distribution of Φ_c and $\tilde{\Phi}_c$ for the MNIST data, with the ReLU function, $p = 784$, $T = 128$ and $c_1 = c_2 = \frac{1}{2}$, with $\mathbf{j}_1 = [\mathbf{1}_{T_1}; \mathbf{0}_{T_2}]$ and $\mathbf{j}_2 = [\mathbf{0}_{T_1}; \mathbf{1}_{T_2}]$. Expectation estimated by averaging over 500 realizations of \mathbf{W} .

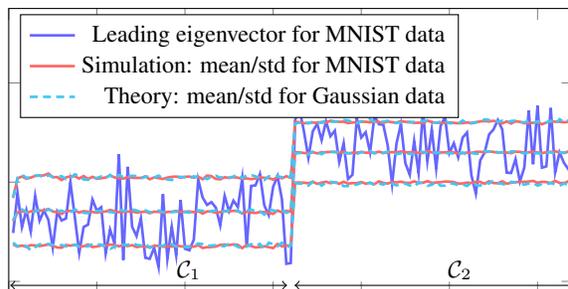


Figure 8. Leading eigenvector of Φ_c for the MNIST and Gaussian mixture data with a width of ± 1 standard deviations (generated from 500 trials) in the settings of Figure 7.

duration. Sets A and B were collected from surface EEG recordings of five healthy volunteers, while sets C, D and E were collected from the EEG records of the pre-surgical diagnosis of five epileptic patients. Here we perform random feature-based spectral clustering on $T = 32, 64$ and 128 randomly picked EEG segments of length $p = 100$ from the dataset. Means and covariances are empirically estimated from the full set (4 097 segments of set B and 4 097 segments of set E). Similar behavior of eigenpairs as for Gaussian mixture models is once more observed in Figure 9 and 10. After k-means classification on the leading two eigenvectors of the (centered) Gram matrix composed of $n = 32$ random features, the accuracies (averaged over 50 runs) are reported in Table 5.

As opposed to the MNIST image recognition task, from Table 5 it is easy to check that the covariance-oriented functions (i.e., $\sigma(t) = |t|$, $\cos(t)$ and $\exp(-t^2/2)$) far outperform any other with almost perfect classification accuracies. It is particularly interesting to note that the popular ReLU function is suboptimal in both tasks, but never performs very badly, thereby offering a good risk-performance tradeoff.

Table 4. Classification accuracies for random feature-based spectral clustering with different $\sigma(t)$ on the MNIST dataset.

	$\sigma(t)$	$T = 32$	$T = 64$	$T = 128$
MEAN-ORIENTED	t	85.31%	88.94%	87.30%
	$\mathbf{1}_{t>0}$	86.00%	82.94%	85.56%
	$\text{sign}(t)$	81.94%	83.34%	85.22%
	$\sin(t)$	85.31%	87.81%	87.50%
	$\text{erf}(t)$	86.50%	87.28%	86.59%
COV-ORIENTED	$ t $	62.81%	60.41%	57.81%
	$\cos(t)$	62.50%	59.56%	57.72%
	$\exp(-\frac{t^2}{2})$	64.00%	60.44%	58.67%
BALANCED	$\text{ReLU}(t)$	82.87%	85.72%	82.27%

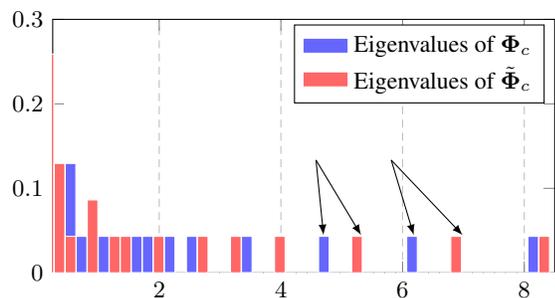
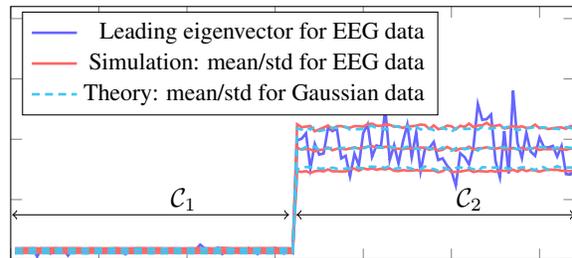

 Figure 9. Eigenvalue distribution of Φ_c and $\tilde{\Phi}_c$ for the epileptic EEG, with the ReLU function, $p = 100$, $T = 128$ and $c_1 = c_2 = \frac{1}{2}$, with $\mathbf{j}_1 = [\mathbf{1}_{T_1}; \mathbf{0}_{T_2}]$ and $\mathbf{j}_2 = [\mathbf{0}_{T_1}; \mathbf{1}_{T_2}]$. Expectation estimated by averaging over 500 realizations of \mathbf{W} .

 Table 5. Classification accuracies for random feature-based spectral clustering with different $\sigma(t)$ on the epileptic EEG dataset.

	$\sigma(t)$	$T = 32$	$T = 64$	$T = 128$
MEAN-ORIENTED	t	71.81%	70.31%	69.58%
	$\mathbf{1}_{t>0}$	65.19%	65.87%	63.47%
	$\text{sign}(t)$	67.13%	64.63%	63.03%
	$\sin(t)$	71.94%	70.34%	68.22%
	$\text{erf}(t)$	69.44%	70.59%	67.70%
COV-ORIENTED	$ t $	99.69%	99.69%	99.50%
	$\cos(t)$	99.00%	99.38%	99.36%
	$\exp(-\frac{t^2}{2})$	99.81%	99.81%	99.77%
BALANCED	$\text{ReLU}(t)$	84.50%	87.91%	90.97%

5. Conclusion

In this article, we have provided a theoretical analysis on random feature-based spectral algorithms for large dimensional data, providing a better understanding of the precise mechanism underlying these methods. Our results show a quite simple relation between the nonlinear function involved in the random feature map (only through two scalars


 Figure 10. Leading eigenvector of Φ_c for the EEG and Gaussian mixture data with a width of ± 1 standard deviation (generated from 500 trials) in the settings of Figure 9.

d_1 and d_2) and the capacity of the latter to discriminate data upon their means and covariances. In obtaining this result, we demonstrated that point-wise nonlinearities can be incorporated into a classical Taylor expansion as a consequence of the concentration phenomenon in high dimensional space. This result was then validated through experimental classification tasks on the MNIST and EEG datasets.

Although Theorem 1 is stated here solely for the functions listed in Table 1, our current line of investigation consists in directly linking the activation function $\sigma(\cdot)$ and the coefficients d_0 , d_1 and d_2 in Table 2 so as to generalize our results and to provide more insights into the attributes of a function that makes it mean- or covariance-oriented; this undertaking is however more technically demanding but still likely achievable through the extension of existing results related to the work of (Cheng & Singer, 2013).

From a point of view of clustering, the crucial information to distinguish different classes is contained in the isolated eigenvalue/eigenvector pairs as shown in (2), the asymptotic behavior of these pairs, as well as their significance for clustering are technically reachable within the analysis framework presented in this paper. When \mathbf{W} follows a non-Gaussian distribution, or when different nonlinearities are combined (e.g., $\cos + \sin$ to get the Gaussian kernel (Rahimi & Recht, 2008)), obtaining the equivalent kernel Φ in the large p, T regime would be a key enabler to gain a deeper understanding under these more elaborate settings.

Besides, this paper can be taken as a first step of the random matrix-based understanding of various learning methods using random features, for example the randomly designed deep neural networks (Lillicrap et al., 2016), the nonlinear activation of which being the main difficulty for a thorough analysis. Moreover, along with recent advances in random matrix analysis (Tiomoko Ali & Couillet, 2016), the hyperparameter d_1/d_2 of utmost importance envisions to be consistently estimated and thus allows for an efficient tuning technique for all nonlinear random feature-based methods.

Acknowledgments

We thank the anonymous reviewers for their comments and constructive suggestions. We would like to acknowledge this work is supported by the ANR Project RMT4GRAPH (ANR-14-CE28-0006) and the Project DeepRMT of La Fondation Supélec.

References

- Andrzejak, R. G., Lehnertz, K., Mormann, F., Rieke, C., David, P., and Elger, C. E. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E*, 64(6):061907, 2001.
- Baik, J., Arous, G. B., Péché, S., et al. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.
- Cheng, X. and Singer, A. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04):1350010, 2013.
- Couillet, R., Benaych-Georges, F., et al. Kernel spectral clustering of large dimensional data. *Electronic Journal of Statistics*, 10(1):1393–1454, 2016.
- El Karoui, N. et al. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010.
- Horn, R. A. and Johnson, C. R. *Matrix analysis*. Cambridge university press, 2012.
- Huang, G.-B., Zhou, H., Ding, X., and Zhang, R. Extreme learning machine for regression and multiclass classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(2):513–529, 2012.
- Keriven, N., Bourrier, A., Gribonval, R., and Pérez, P. Sketching for large-scale learning of mixture models. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pp. 6190–6194. IEEE, 2016.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- LeCun, Y., Cortes, C., and Burges, C. J. The MNIST database of handwritten digits, 1998.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7, 2016.
- Louart, C., Liao, Z., Couillet, R., et al. A random matrix approach to neural networks. *The Annals of Applied Probability*, 28(2):1190–1248, 2018.
- Maas, A. L., Hannun, A. Y., and Ng, A. Y. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, 2013.
- Marčenko, V. A. and Pastur, L. A. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457, 1967.
- Ng, A. Y., Jordan, M. I., and Weiss, Y. On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pp. 849–856, 2002.
- Pennington, J. and Worah, P. Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, pp. 2634–2643, 2017.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pp. 1177–1184, 2008.
- Scardapane, S. and Wang, D. Randomness in neural networks: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 7(2), 2017.
- Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Schölkopf, B. and Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- Silverstein, J. W. and Bai, Z. On the empirical distribution of eigenvalues of a class of large dimensional random matrices. *Journal of Multivariate analysis*, 54(2):175–192, 1995.
- Tiomoko Ali, H. and Couillet, R. Spectral community detection in heterogeneous large networks. *arXiv preprint arXiv:1611.01096*, 2016.
- Vedaldi, A. and Zisserman, A. Efficient additive kernels via explicit feature maps. *IEEE transactions on pattern analysis and machine intelligence*, 34(3):480–492, 2012.
- Williams, C. K. Computing with infinite networks. *Advances in neural information processing systems*, pp. 295–301, 1997.

Supplementary Material

On the Spectrum of Random Features Maps of High Dimensional Data

A. Proof of Theorem 1

To obtain the result presented in Theorem 1, we begin by recalling the expression of the average kernel matrix Φ , with

$$\Phi_{i,j} = \Phi(\mathbf{x}_i, \mathbf{x}_j) = \mathbb{E}_{\mathbf{w}} \mathbf{G}_{ij} = \mathbb{E}_{\mathbf{w}} \sigma(\mathbf{w}^\top \mathbf{x}_i) \sigma(\mathbf{w}^\top \mathbf{x}_j).$$

For $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$, one resorts to the integral calculus for standard Gaussian distribution in \mathbb{R}^p , which can be further reduced to a double integral as shown in (Williams, 1997; Louart et al., 2018) and results in the expressions in Table 1.

Then, from the discussion in Section 2, we have the following expansions for $\mathbf{x}_i \in \mathcal{C}_a$, $\mathbf{x}_j \in \mathcal{C}_b$, with $i \neq j$,

$$\mathbf{x}_i^\top \mathbf{x}_j = \underbrace{\omega_i^\top \omega_j}_{O(p^{-1/2})} + \underbrace{\mu_a^\top \mu_b / p + \mu_a^\top \omega_j / \sqrt{p} + \mu_b^\top \omega_i / \sqrt{p}}_{O(p^{-1})} \quad (3)$$

$$\|\mathbf{x}_i\|^2 = \underbrace{\tau}_{O(1)} + \underbrace{\text{tr}(\mathbf{C}_a^\circ) / p + \|\omega_i\|^2 - \text{tr}(\mathbf{C}_a) / p}_{O(p^{-1/2})} + \underbrace{\|\mu_a\|^2 / p + 2\mu_a^\top \omega_i / \sqrt{p}}_{O(p^{-1})} \quad (4)$$

which further allows one to *linearize* the nonlinear function of $\mathbf{x}_i, \mathbf{x}_j$ in Table 1 via a Taylor expansion to obtain an *entry-wise* approximation of the key matrix Φ .

Nonetheless, this entry-wise approximation does not ensure a vanishing difference in terms of operator norm in the large p, T regime under consideration. Taking the popular Marčenko–Pastur law (Marčenko & Pastur, 1967) for example: consider a random matrix $\mathbf{W} \in \mathbb{R}^{n \times p}$ with i.i.d. standard Gaussian entries. Then, as $n, p \rightarrow \infty$ with $\frac{n}{p} \rightarrow c \in (0, \infty)$, entry-wisely we have that the entry (i, j) of the matrix $\frac{1}{p} \mathbf{W} \mathbf{W}^\top$ converge to 1 if $i = j$ and 0 otherwise, meaning that the sample covariance matrix $\frac{1}{p} \mathbf{W} \mathbf{W}^\top$ seemingly “converge to” an identity matrix (which is indeed the true population covariance). But it is well known (Marčenko & Pastur, 1967) that the eigenvalue distribution of $\frac{1}{p} \mathbf{W} \mathbf{W}^\top$ converges (almost surely so) to a continuous measure (the popular Marčenko–Pastur distribution) compactly supported on $[(1 - \sqrt{c})^2, (1 + \sqrt{c})^2]$, which is evidently different from the eigenvalue distribution $\delta_{x=1}$ of \mathbf{I}_n .

As a consequence, a more careful control of the entry-wise expansion of Φ must be performed to ensure a vanishing expansion error in terms of operator norm. To this end, we follow the previous work of (El Karoui et al., 2010; Couillet et al., 2016) and consider the full matrix contribution.

In the following we proceed the aforementioned manipulations on the ReLU function as an example, derivations of other functions follow the same procedure and are thus omitted.

Proof of $\sigma(t) = \text{ReLU}(t)$. We start with the computation of $\Phi(\mathbf{a}, \mathbf{b})$. For $\sigma(t) = \text{ReLU}(t)$, with the classical Gram-Schmidt process we obtain

$$\begin{aligned} \Phi(\mathbf{a}, \mathbf{b}) &= \mathbb{E}_{\mathbf{w}} \sigma(\mathbf{w}^\top \mathbf{a}) \sigma(\mathbf{w}^\top \mathbf{b}) = (2\pi)^{-\frac{p}{2}} \int_{\mathbb{R}^p} \sigma(\mathbf{w}^\top \mathbf{a}) \sigma(\mathbf{w}^\top \mathbf{b}) e^{-\frac{1}{2} \|\mathbf{w}\|^2} d\mathbf{w} \\ &= \frac{1}{2\pi} \int_{\mathbb{R}} \int_{\mathbb{R}} \sigma(\tilde{w}_1 \tilde{a}_1) \sigma(\tilde{w}_1 \tilde{b}_1 + \tilde{w}_2 \tilde{b}_2) e^{-\frac{1}{2} (\tilde{w}_1^2 + \tilde{w}_2^2)} d\tilde{w}_1 d\tilde{w}_2 \\ &= \frac{1}{2\pi} \int_{\mathbb{R}^2} \sigma(\tilde{\mathbf{w}}^\top \tilde{\mathbf{a}}) \sigma(\tilde{\mathbf{w}}^\top \tilde{\mathbf{b}}) e^{-\frac{1}{2} \|\tilde{\mathbf{w}}\|^2} d\tilde{\mathbf{w}} \\ &= \frac{1}{2\pi} \int_{\min(\tilde{\mathbf{w}}^\top \tilde{\mathbf{a}}, \tilde{\mathbf{w}}^\top \tilde{\mathbf{b}}) \geq 0} \tilde{\mathbf{w}}^\top \tilde{\mathbf{a}} \cdot \tilde{\mathbf{w}}^\top \tilde{\mathbf{b}} \cdot e^{-\frac{1}{2} \|\tilde{\mathbf{w}}\|^2} d\tilde{\mathbf{w}} \end{aligned}$$

where $\tilde{a}_1 = \|\mathbf{a}\|$, $\tilde{b}_1 = \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\|}$, $\tilde{b}_2 = \|\mathbf{b}\| \sqrt{1 - \frac{(\mathbf{a}^\top \mathbf{b})^2}{\|\mathbf{a}\|^2 \|\mathbf{b}\|^2}}$ and we denote $\tilde{\mathbf{w}} = [\tilde{w}_1, \tilde{w}_2]^\top$, $\tilde{\mathbf{a}} = [\tilde{a}_1, 0]^\top$ and $\tilde{\mathbf{b}} = [\tilde{b}_1, \tilde{b}_2]^\top$.

With a simple geometric representation we observe

$$\{\tilde{\mathbf{w}} \mid \min(\tilde{\mathbf{w}}^\top \tilde{\mathbf{a}}, \tilde{\mathbf{w}}^\top \tilde{\mathbf{b}}) \geq 0\} = \left\{ r \cos(\theta) + r \sin(\theta) \mid r \geq 0, \theta \in \left[\theta_0 - \frac{\pi}{2}, \frac{\pi}{2} \right] \right\}$$

with $\theta_0 \equiv \arccos\left(\frac{\tilde{b}_1}{\|\tilde{\mathbf{b}}\|}\right) = \frac{\pi}{2} - \arcsin\left(\frac{\tilde{b}_1}{\|\tilde{\mathbf{b}}\|}\right)$. Therefore with a polar coordinate change of variable we deduce, for $\sigma(t) = \text{ReLU}(t)$ that

$$\begin{aligned} \Phi(\mathbf{a}, \mathbf{b}) &= \tilde{a}_1 \frac{1}{2\pi} \int_{\theta_0 - \frac{\pi}{2}}^{\frac{\pi}{2}} \cos(\theta) \left(\tilde{b}_1 \cos(\theta) + \tilde{b}_2 \sin(\theta) \right) d\theta \int_{\mathbb{R}^+} r^3 e^{-\frac{1}{2}r^2} dr \\ &= \frac{1}{2\pi} \|\mathbf{a}\| \|\mathbf{b}\| \left(\sqrt{1 - \angle(\mathbf{a}, \mathbf{b})^2} + \angle(\mathbf{a}, \mathbf{b}) \arccos(-\angle(\mathbf{a}, \mathbf{b})) \right) \end{aligned}$$

with $\angle(\mathbf{a}, \mathbf{b}) \equiv \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$ as Table 1.

For a second step, with the expressions in (3) and (4) we perform a Taylor expansion to get

$$\angle(\mathbf{x}_i, \mathbf{x}_j) = \underbrace{\frac{1}{\tau} \boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j}_{O(p^{-1/2})} + \underbrace{\frac{1}{\tau} \left(\frac{1}{p} \boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b + \frac{1}{\sqrt{p}} \boldsymbol{\mu}_a^\top \boldsymbol{\omega}_j + \frac{1}{\sqrt{p}} \boldsymbol{\mu}_b^\top \boldsymbol{\omega}_i \right) - \frac{\boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j}{2\tau^2} \left(\frac{1}{p} \text{tr} \mathbf{C}_a^\circ + \phi_i + \frac{1}{p} \text{tr} \mathbf{C}_b^\circ + \phi_j \right)}_{O(p^{-1})} + O(p^{-\frac{3}{2}})$$

where we recall $\phi_i = \|\boldsymbol{\omega}_i\|^2 - \mathbb{E}\|\boldsymbol{\omega}_i\|^2 = \|\boldsymbol{\omega}_i\|^2 - \text{tr}(\mathbf{C}_a)/p$ that is of order $O(p^{-1/2})$. Note that the third term $\boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j$ ($\text{tr} \mathbf{C}_a^\circ/p + \phi_i + \text{tr} \mathbf{C}_b^\circ/p + \phi_j$), being of order $O(p^{-1})$, gives rise to a matrix of vanishing operator norm (Couillet et al., 2016), we thus conclude by stating that

$$\angle(\mathbf{x}_i, \mathbf{x}_j) = \underbrace{\frac{1}{\tau} \boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j}_{O(p^{-1/2})} + \underbrace{\frac{1}{\tau} \left(\frac{1}{p} \boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b + \frac{1}{\sqrt{p}} \boldsymbol{\mu}_a^\top \boldsymbol{\omega}_j + \frac{1}{\sqrt{p}} \boldsymbol{\mu}_b^\top \boldsymbol{\omega}_i \right)}_{O(p^{-1})} + O(p^{-\frac{3}{2}}).$$

Since $\angle(\mathbf{x}_i, \mathbf{x}_j) = O(p^{-\frac{1}{2}})$ we sequentially Taylor-expand $\sqrt{1 - \angle(\mathbf{x}_i, \mathbf{x}_j)^2}$ and $\arccos(-\angle(\mathbf{x}_i, \mathbf{x}_j))$ to the order of $O(p^{-\frac{3}{2}})$ as

$$\begin{aligned} \sqrt{1 - \angle(\mathbf{x}_i, \mathbf{x}_j)^2} &= 1 - \frac{1}{2} \angle(\mathbf{x}_i, \mathbf{x}_j)^2 + O(p^{-\frac{3}{2}}) \\ \arccos(-\angle(\mathbf{x}_i, \mathbf{x}_j)) &= \frac{\pi}{2} + \angle(\mathbf{x}_i, \mathbf{x}_j) + O(p^{-\frac{3}{2}}). \end{aligned}$$

As such, we conclude for $\sigma(t) = \text{ReLU}(t)$ that

$$\begin{aligned} \Phi(\mathbf{x}_i, \mathbf{x}_j) &= \frac{1}{2\pi} \|\mathbf{x}_i\| \|\mathbf{x}_j\| \left(\sqrt{1 - \angle(\mathbf{x}_i, \mathbf{x}_j)^2} + \angle(\mathbf{x}_i, \mathbf{x}_j) \arccos(-\angle(\mathbf{x}_i, \mathbf{x}_j)) \right) \\ &= \frac{1}{2\pi} \|\mathbf{x}_i\| \|\mathbf{x}_j\| \left(1 + \frac{\pi}{2} \angle(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{2} \angle(\mathbf{x}_i, \mathbf{x}_j)^2 \right) + O(p^{-\frac{3}{2}}) \\ &= \frac{1}{4} \mathbf{x}_i^\top \mathbf{x}_j + \frac{1}{2\pi} \|\mathbf{x}_i\| \|\mathbf{x}_j\| + \frac{1}{4\pi} \mathbf{x}_i^\top \mathbf{x}_j \angle(\mathbf{x}_i, \mathbf{x}_j) + O(p^{-\frac{3}{2}}). \end{aligned}$$

Consequently we get the generic form for all functions $\sigma(\cdot)$ listed in Table 1 is given by

$$\begin{aligned} \Phi_{i,j} \equiv \Phi(\mathbf{x}_i, \mathbf{x}_j) &= \underbrace{c_0}_{O(1)} + \underbrace{c_1 (\mathbf{t}_a/\sqrt{p} + \phi_i + \mathbf{t}_b/\sqrt{p} + \phi_j) + c_2 \boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j}_{O(p^{-1/2})} \\ &+ \underbrace{c_3 (\mathbf{t}_a/\sqrt{p} + \phi_i) (\mathbf{t}_b/\sqrt{p} + \phi_j) + c_4 ((\mathbf{t}_a/\sqrt{p} + \phi_i)^2 + (\mathbf{t}_b/\sqrt{p} + \phi_j)^2) + c_5 (\boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j)^2}_{O(p^{-1})} \\ &+ \underbrace{c_7 ((\|\boldsymbol{\mu}_a\|^2 + \|\boldsymbol{\mu}_b\|^2)/p + 2(\boldsymbol{\mu}_a^\top \boldsymbol{\omega}_i + \boldsymbol{\mu}_b^\top \boldsymbol{\omega}_j)/\sqrt{p}) + c_8 (\boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b/p + \boldsymbol{\mu}_a^\top \boldsymbol{\omega}_j/\sqrt{p} + \boldsymbol{\mu}_b^\top \boldsymbol{\omega}_i/\sqrt{p})}_{O(p^{-1})} + O(p^{-\frac{3}{2}}) \end{aligned}$$

where we recall the definition $\mathbf{t} \equiv \{\text{tr} \mathbf{C}_a^\circ / \sqrt{p}\}_{a=1}^K$. In particular, we have for the ReLU nonlinearity $c_0 = \frac{\tau}{2\pi}$, $c_1 = c_7 = \frac{1}{4\pi}$, $c_2 = c_8 = \frac{1}{4}$, $c_3 = \frac{1}{8\pi\tau}$, $c_4 = -\frac{1}{16\pi\tau}$, $c_5 = \frac{1}{4\pi\tau}$ and $c_6 = 0$.

We then observe that, for all functions $\sigma(\cdot)$ listed in Table 1, we have $c_7 = c_1$ and $c_8 = c_2$. Besides, using the fact that

$$(\boldsymbol{\omega}_i^\top \boldsymbol{\omega}_j)^2 = \text{tr}(\mathbf{C}_a \mathbf{C}_b) / p^2 + O(p^{-\frac{3}{2}})$$

and considering also the diagonal terms (with $i = j$) by adding the term $c_9 \mathbf{I}_T$, we finally get

$$\begin{aligned} \Phi &= c_0 \mathbf{1}_T \mathbf{1}_T^\top + c_1 \left(\phi \mathbf{1}_T^\top + \mathbf{1}_T \phi^\top + \left\{ \frac{\mathbf{t}_a \mathbf{1}_{T_a}}{\sqrt{p}} \right\}_{a=1}^K \mathbf{1}_T^\top + \mathbf{1}_T \left\{ \frac{\mathbf{t}_b \mathbf{1}_{T_b}^\top}{\sqrt{p}} \right\}_{b=1}^K \right) + c_2 \boldsymbol{\Omega}^\top \boldsymbol{\Omega} \\ &+ c_3 \left(\phi \phi^\top + \phi \left\{ \frac{\mathbf{t}_b \mathbf{1}_{T_b}^\top}{\sqrt{p}} \right\}_{b=1}^K + \left\{ \frac{\mathbf{t}_a \mathbf{1}_{T_a}}{\sqrt{p}} \right\}_{a=1}^K \phi^\top + \left\{ \mathbf{t}_a \mathbf{t}_b \frac{\mathbf{1}_{T_a} \mathbf{1}_{T_b}^\top}{p} \right\}_{a,b=1}^K \right) \\ &+ c_4 \left((\phi^2) \mathbf{1}_T^\top + \mathbf{1}_T (\phi^2)^\top + 2 \left(\mathcal{D} \left\{ \mathbf{t}_a \mathbf{1}_{T_a} \right\}_{a=1}^K \phi \frac{\mathbf{1}_T^\top}{\sqrt{p}} \right) + 2 \left(\frac{\mathbf{1}_T}{\sqrt{p}} \phi^\top \mathcal{D} \left\{ \mathbf{t}_b \mathbf{1}_{T_b} \right\}_{b=1}^K \right) \right. \\ &+ \left. \left\{ \mathbf{t}_a^2 \frac{\mathbf{1}_{T_a}}{p} \right\}_{a=1}^K \mathbf{1}_T^\top + \mathbf{1}_T \left\{ \mathbf{t}_b^2 \frac{\mathbf{1}_{T_b}^\top}{p} \right\}_{b=1}^K \right) + c_5 \left\{ \text{tr}(\mathbf{C}_a \mathbf{C}_b) \frac{\mathbf{1}_{T_a} \mathbf{1}_{T_b}^\top}{p^2} \right\}_{a,b=1}^K \\ &+ c_1 \left(\left\{ \|\boldsymbol{\mu}_a\|^2 \frac{\mathbf{1}_{T_a}}{p} \right\}_{a=1}^K \mathbf{1}_T^\top + \mathbf{1}_T \left\{ \|\boldsymbol{\mu}_b\|^2 \frac{\mathbf{1}_{T_b}^\top}{p} \right\}_{b=1}^K + \frac{2}{\sqrt{p}} \left\{ \boldsymbol{\Omega}_a^\top \boldsymbol{\mu}_a \mathbf{1}_{T_a}^\top \right\}_{a=1}^K + \frac{2}{\sqrt{p}} \left\{ \mathbf{1}_{T_b} \boldsymbol{\mu}_b^\top \boldsymbol{\Omega}_b \right\}_{b=1}^K \right) \\ &+ c_2 \left(\frac{1}{p} \left\{ \mathbf{1}_{T_a} \boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b \mathbf{1}_{T_b}^\top \right\}_{a,b=1}^K + \frac{1}{\sqrt{p}} \left\{ \boldsymbol{\Omega}_a^\top \boldsymbol{\mu}_b \mathbf{1}_{T_b}^\top + \mathbf{1}_{T_a} \boldsymbol{\mu}_a^\top \boldsymbol{\Omega}_b \right\}_{a,b=1}^K \right) + c_9 \mathbf{I}_T + O_{\|\cdot\|}(p^{-\frac{1}{2}}) \end{aligned}$$

where we denote $\phi^2 \equiv [\phi_1^2, \dots, \phi_T^2]^\top$ and $O_{\|\cdot\|}(p^{-\frac{1}{2}})$ a matrix of operator norm of order $O(p^{-1/2})$ as $p \rightarrow \infty$.

Recalling that for $\mathbf{P} \equiv \mathbf{I}_T - \frac{1}{T} \mathbf{1}_T \mathbf{1}_T^\top$, we have $\mathbf{P} \mathbf{1}_T = \mathbf{1}_T \mathbf{P} = \mathbf{0}$ and therefore

$$\begin{aligned} \Phi_c &\equiv \mathbf{P} \Phi \mathbf{P} = c_2 \mathbf{P} \boldsymbol{\Omega}^\top \boldsymbol{\Omega} \mathbf{P} + c_3 \mathbf{P} \left(\phi \phi^\top + \phi \left\{ \frac{\mathbf{t}_b \mathbf{1}_{T_b}^\top}{\sqrt{p}} \right\}_{b=1}^K + \left\{ \frac{\mathbf{t}_a \mathbf{1}_{T_a}}{\sqrt{p}} \right\}_{a=1}^K \phi^\top + \left\{ \mathbf{t}_a \mathbf{t}_b \frac{\mathbf{1}_{T_a} \mathbf{1}_{T_b}^\top}{p} \right\}_{a,b=1}^K \right) \mathbf{P} \\ &+ c_5 \mathbf{P} \left\{ \text{tr}(\mathbf{C}_a \mathbf{C}_b) \frac{\mathbf{1}_{T_a} \mathbf{1}_{T_b}^\top}{p^2} \right\}_{a,b=1}^K \mathbf{P} + c_2 \mathbf{P} \left(\frac{1}{p} \left\{ \mathbf{1}_{T_a} \boldsymbol{\mu}_a^\top \boldsymbol{\mu}_b \mathbf{1}_{T_b}^\top \right\}_{a,b=1}^K + \frac{1}{\sqrt{p}} \left\{ \boldsymbol{\Omega}_a^\top \boldsymbol{\mu}_b \mathbf{1}_{T_b}^\top + \mathbf{1}_{T_a} \boldsymbol{\mu}_a^\top \boldsymbol{\Omega}_b \right\}_{a,b=1}^K \right) \mathbf{P} \\ &+ c_7 \mathbf{P} + O_{\|\cdot\|}(p^{-1/2}) \equiv \mathbf{P} \tilde{\Phi} \mathbf{P} + O_{\|\cdot\|}(p^{-1/2}). \end{aligned}$$

We further observe that, for all functions $\sigma(\cdot)$ listed in Table 1 we have $c_5 = 2c_3$ and let $d_0 = c_7$, $d_1 = c_2$, $d_2 = c_3 = c_5/2$ we obtain the expression of $\tilde{\Phi}$ in Theorem 1, which concludes the proof. \square