



HAL
open science

Feature uncertainty bounds for explicit feature maps and large robust nonlinear SVM classifiers

Nicolas Couellan, Sophie Jan

► **To cite this version:**

Nicolas Couellan, Sophie Jan. Feature uncertainty bounds for explicit feature maps and large robust nonlinear SVM classifiers. *Annals of Mathematics and Artificial Intelligence*, 2020, 88 (1-3), pp.269-289. 10.1007/s10472-019-09676-0 . hal-01718729

HAL Id: hal-01718729

<https://hal.science/hal-01718729v1>

Submitted on 27 Feb 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Feature uncertainty bounding schemes for large robust nonlinear SVM classifiers

Nicolas Couellan* Sophie Jan†

Abstract

We consider the binary classification problem when data are large and subject to unknown but bounded uncertainties. We address the problem by formulating the nonlinear support vector machine training problem with robust optimization. To do so, we analyze and propose two bounding schemes for uncertainties associated to random approximate features in low dimensional spaces. The proposed techniques are based on Random Fourier Features and the Nyström methods. The resulting formulations can be solved with efficient stochastic approximation techniques such as stochastic (sub)-gradient, stochastic proximal gradient techniques or their variants.

1 Introduction

When designing classification models, the common idea is to consider that training samples are certain and not subject to noise. In reality, this assumption may be far from being verified. There are several sources of possible noise. For example, data may be noisy by nature. In biology, for instance, measurements on living cells under the microscope are unprecise due to the movement of the studied organism. Furthermore, any phenomenon that occurs at a smaller scale than the given microscope resolution is measured with errors. Uncertainty is sometimes also introduced as a way to handle missing data. If some of the sample attributes are not known but can be estimated, they can be replaced by their estimation, assuming some level of uncertainty. In the work that is presented here, we are interested in constructing optimization-based classification models that can account for uncertainty in the data and provide some robustness to noise. To do so, some knowledge on the noise is required. Most of the time, assuming that we know the probability distribution of the noisy perturbations is unrealistic. However, depending on the context, one may have knowledge on the mean and variance or simply the magnitude of the perturbations. The

*Institut de Mathématiques de Toulouse, Université Paul Sabatier, 118 route de Narbonne 31062 Toulouse Cedex 9, France - nicolas.couellan@math.univ-toulouse.fr

†Institut de Mathématiques de Toulouse, Université Paul Sabatier, 118 route de Narbonne 31062 Toulouse Cedex 9, France

quantity and the nature of information available will be critical for the design of models that are immune to noise.

In binary classification problems, the idea is to compute a decision rule based on empirical data in order to predict the class on new observations. If empirical data is uncertain and no robustness is incorporated in the model, the decision rule will also lead to uncertain decision. In the context of Support Vector Machine (SVM) classification [22], there are several ways to construct robust counterparts of the classification problem, see for example [2]. In these techniques, one considers that data points can be anywhere within an uncertainty set centered at a nominal point. The uncertainty set may have various shapes and dimension depending on the application and the level of noise. In the chance constrained SVM method, instead of ensuring that as many possible points are on the right side of the separating hyperplane, the idea is to ensure that a large probability of them will be correctly classified [3, 2]. Alternatively, one may want to ensure that no matter where the data points are in the uncertainty set, the separating hyperplane will ensure that it will always be on the correct side. This is known as the worst case robust SVM approach [19, 24]. In both methods, one can show that the resulting classification problem can be formulated as a Second Order Cone Program (SOCP).

Classification problems are usually very large optimization problems due to the great number of samples available. Furthermore, so far, to the best of our knowledge, robust SVM formulations are only dealing with linear classification. If uncertainties are large, additional non-separability is introduced in the data that may limit the predictive performance of robust SVM models.

To propose solutions to these challenges, we present several new models and SVM classification approaches. We first describe the concept of robust Support Vector Machine (SVM) classification and its computation through stochastic optimization techniques when the number of samples is large. We further explain the concept of safe SVM we have introduced to overcome the conservatism issues in classical robust optimization techniques. We then develop new extensions to robust nonlinear classification models where we propose computational schemes to bound uncertainties in the feature space.

2 Notations

We define the following notations:

- $\Sigma^{1/2}$ the matrix in $\mathbb{R}^{n \times n}$ associated to the positive definite matrix Σ such that $\Sigma^{1/2}(\Sigma^{1/2})^\top = \Sigma$;
- $\Sigma^{-1/2}$ the matrix $(\Sigma^{1/2})^{-1}$;
- $\Sigma^{\top/2}$ the matrix $(\Sigma^{1/2})^\top$;
- $\mathcal{N}(\mu, \nu^2)$ the normal distribution with mean μ and variance ν^2 .

3 Robust SVM classification models: preliminaries

We first recall the standard SVM methodology [22, 14, 8] to find the maximum margin separating hyperplane between two classes of data points.

Let $x_i \in \mathbb{R}^n$ be a collection of input training vectors for $i = 1, \dots, L$ and $y_i \in \{-1, 1\}$ be their corresponding labels. The SVM classification problem to separate the training vectors into the -1 and 1 classes can be formulated as follows:

$$\begin{aligned} \min_{w, b, \xi} \quad & \frac{\lambda}{2} \|w\|^2 + \sum_{i=1}^L \xi_i & (1) \\ \text{subject to} \quad & y_i(\langle w, x_i \rangle + b) + \xi_i \geq 1 \quad i = 1, \dots, L \\ & \xi_i \geq 0 \quad i = 1, \dots, L \end{aligned}$$

where ξ_i are slack variables added to the constraints and in the objective as a penalty term, λ is a problem specific constant controlling the trade-off between margin (generalization) and classification, and $\langle w, x \rangle + b = 0$ defines the separating hyperplane with w the normal vector to the separating hyperplane and b its relative position to the origin.

Consider now instead a set of noisy training vectors $\{\tilde{x}_i \in \mathbb{R}^n, i = 1, \dots, L\}$ where $\tilde{x}_i = x_i + \Delta x_i$ for all i in $\{1, \dots, L\}$ and Δx_i is a random perturbation, Problem (1) becomes:

$$\begin{aligned} \min_{w, b, \xi} \quad & \frac{\lambda}{2} \|w\|^2 + \sum_{i=1}^L \xi_i & (2) \\ \text{subject to} \quad & y_i(\langle w, x_i + \Delta x_i \rangle + b) + \xi_i \geq 1, \quad \forall i = 1, \dots, L \\ & \xi_i \geq 0, \quad \forall i = 1, \dots, L. \end{aligned}$$

Observe that the first constraint involves the random variable Δx_i and Problem (2) cannot be solved as such. Extra knowledge on the perturbations is needed to transform (2) into a deterministic and numerically solvable problem. If the uncertainty set is bounded and the bound is known, one can consider that

$$\|\Sigma_i^{-1/2} \Delta x_i\|_p \leq \gamma_i, \quad \text{with } \gamma_i > 0, \forall i \in \{1, \dots, L\} \text{ and } p \geq 1.$$

Various choices of Σ_i and p will lead to various types of uncertainties such as for example box-shaped uncertainty ($\|\tilde{x}_i - x_i\|_\infty \leq \gamma_i$), spherical uncertainty ($\|\tilde{x}_i - x_i\|_2 \leq \gamma_i$) or ellipsoidal uncertainty ($(\tilde{x}_i - x_i)^\top \Sigma_i^{-1} (\tilde{x}_i - x_i) \leq \gamma_i^2$ for some positive definite matrix Σ_i).

To design a robust model, one has to satisfy the linear inequality constraint in Problem (2) for every realizations of Δx_i . This can be done by ensuring the constraint in the worst case scenario for Δx_i , leading to the following robust

counterpart optimization problem:

$$\begin{aligned}
& \min_{w,b,\xi} && \frac{\lambda}{2}\|w\|^2 + \sum_{i=1}^L \xi_i && (3) \\
\text{subject to} &&& \min_{\|\Sigma_i^{-1/2}\Delta x_i\|_p \leq \gamma_i} y_i(\langle w, x_i + \Delta x_i \rangle + b) + \xi_i \geq 1, && \forall i = 1, \dots, L, \\
&&& \xi_i \geq 0, && \forall i = 1, \dots, L.
\end{aligned}$$

Denoting $\|\cdot\|_q$ the dual norm of $\|\cdot\|_p$, we have

$$\min_{\|\Sigma_i^{-1/2}\Delta x_i\|_p \leq \gamma_i} y_i \langle w, \Delta x_i \rangle = -\gamma_i \left\| \Sigma_i^{\top/2} w \right\|_q.$$

Therefore, Problem (3) can be rewritten as

$$\begin{aligned}
& \min_{w,b,\xi} && \frac{\lambda}{2}\|w\|^2 + \sum_{i=1}^L \xi_i && (4) \\
\text{subject to} &&& y_i(\langle w, x_i \rangle + b) + \xi_i - \gamma_i \left\| \Sigma_i^{\top/2} w \right\|_q \geq 1, && \forall i = 1, \dots, L, \\
&&& \xi_i \geq 0, && \forall i = 1, \dots, L.
\end{aligned}$$

This problem is a Second Order Cone Program (SOCP). There has been previous investigation on this formulation where the problem was solved using a primal-dual IPM [1] and solver packages such as [17, 12]. The complexity of primal-dual IPM is $O(s^3)$ operations per iteration and the number of iterations is $O(\sqrt{s} \times \log(1/\epsilon))$ where ϵ is the required precision and s the number of cones in the problem. In Problem (4), the number of cones is $L + 1$ (L is the size of the dataset). Therefore these methods do not scale well as the size of the dataset increases. However, for small and medium size problems, they have given satisfying results [20, 19]. For large problems, we propose alternative techniques in the next section.

Chance constrained approach

Observe that if one has additional knowledge on the bound of the perturbations (ex: mean and variance of the noise distribution), one may want to consider the chance constrained approach rather than the worst case methodology above. Doing so, we can replace in Problem (2) the constraint $y_i(\langle w, x_i + \Delta x_i \rangle + b) + \xi_i \geq 1$ by

$$\mathbb{P}_{\Delta x_i} [y_i(\langle w, x_i + \Delta x_i \rangle + b) + \xi_i \geq 1] \geq 1 - \epsilon_i, \quad \forall i \in \{1, \dots, L\},$$

where ϵ_i is a given positive confidence level. In other words, instead of ensuring the constraint for all realizations of Δx_i , we would like to ensure its probability of being satisfied to be greater than some threshold value $1 - \epsilon_i$. Indeed, it was shown in [16] that if Δx_i follows a normal distribution $\mathcal{N}(0, \Sigma_i)$, the chance

constraint problem can be formulated exactly as Problem (4) if p is taken as 2 and ϵ_i such that $\gamma_i = \sqrt{\epsilon_i/(1-\epsilon_i)}$. The proof makes use of the multivariate Chebyshev inequality to replace the probability of ensuring the constraint by the Chebyshev bound. The two approaches are therefore related when $p = 2$. In the work that is described next, we have chosen to exploit the worst case model as less assumptions are taken on the uncertainty set since p can take any positive value greater than 1 and we do not require explicit knowledge of variance of the perturbations.

Large scale robust framework

We now consider the case where the dataset is very large and SOCP techniques cannot be used in practice. We reformulate the problem (4) as an equivalent unconstrained optimization problem:

$$\min_{w,b} f(w,b) = \frac{\lambda}{2}\|w\|^2 + \sum_{i=1}^L \ell_{\Sigma_i}(y_i, \langle w, x_i \rangle + b), \quad (5)$$

where ℓ_{Σ_i} is the robust loss function defined by

$$\ell_{\Sigma_i}(y_i, \langle w, x_i \rangle + b) = \max \left\{ 0, 1 - y_i(\langle w, x_i \rangle + b) + \gamma_i \|\Sigma_i^{\top/2} w\|_q \right\}. \quad (6)$$

There are many techniques available to solve problem (5). When L is very large, stochastic approximation techniques may be considered as they have very low per-iteration complexity. For example, stochastic proximal methods have proven to be very effective for this type of large decomposable problems [23]. The idea is to split the objective f into two components $G(w,b) = g(w,b) + h(w)$ where

$$g(w,b) = \sum_{l=1}^L g_l(w,b), \quad g_l(w,b) = \ell_{\Sigma_l}(y_l, \langle w, x_l \rangle + b), \quad h(w) = \frac{\lambda}{2}\|w\|^2,$$

to perform a stochastic forward step on the decomposable component g by taking a gradient approximation using only one random term g_l of the decomposable sum, and finally to carry out a backward proximal step using h .

Controlling the robustness of the model

Chance constraints and worst case robust formulations are based on a priori information on the noise bound or its variance. If these estimations are erroneous and over-estimated for example, the models may be too conservative. Additionally, worst case situations may be too pessimistic and may introduce additional non separability in the problem resulting in lower generalization performance of the model. This is the reason why we have proposed an alternative adaptive robust model that will have the advantage of being less conservative without taking any additional assumption on the probability distribution [7]. The idea

is to consider an adjustable subset of the support of the probability distribution of the uncertainties. Assuming an ellipsoidal uncertainty set, a reduced uncertainty set (reduced ellipsoid) is computed so as to minimize a generalization error. The model is referred to as *safe SVM* rather than robust SVM. Mathematically, this is achieved by introducing a variable σ defining a Σ_σ matrix of a reduced ellipsoid where σ is the vector of lengths of the ellipsoid along its axes. The resulting model can be cast as a bi-level program as follows:

$$\begin{aligned}
\min_{\sigma \in \mathbb{R}^n} \quad & \sum_{j=1}^N \ell(y_j^v, \langle w^*, x_j^v \rangle + b^*) \\
\text{s.t.} \quad & \sigma_{min} \leq \sigma_t \leq \sigma_{max} \quad \forall t = 1, \dots, n \\
& (w^*, b^*) = \operatorname{argmin}_{w, b} \frac{\lambda}{2} \|w\|^2 + \sum_{i=1}^L \ell_{\Sigma_{\sigma_t}}(y_i, \langle w, x_i \rangle + b)
\end{aligned} \tag{7}$$

where (x_j^v, y_j^v) , for $j = 1, \dots, N$, are the sample vectors and their labels from the validation fold. The upper and lower bounds σ_{max} and σ_{min} are parameters that control the minimum and maximum amount of uncertainty we would like to take into account in the model. They can be taken as rough bounds of the perturbations. The function ℓ is the standard SVM hinge loss function that will estimate the generalization error for one test sample and ℓ_{Σ_σ} is the robust loss function as defined in (6). The matrix Σ_σ is defined as $(\operatorname{diag}(\sigma))^2$.

In [7], we have solved Problem (7) with a stochastic bi-level gradient algorithm [6]. Using a *robust error* measure as an indicator of performance of the robustness of the model, we have shown that the proposed model achieves a better *robust error* on several public datasets when compared to the SOCP formulation or formulation (5). Using gaussian noise, we have also shown that the volume of the reduced ellipsoid computed by the technique was indeed increasing with variance of the noise, confirming that the model was auto adjusting its robustness to the noise in the data.

4 Nonlinear robust SVM classification models

There are situations where the model (1) will not be able to compute a sufficiently good linear separation between the classes if the data are highly nonlinearly separable. This phenomenon is actually even more present when noise is taken into account since robust models introduce additional non separability in the data. For this reason, it is interesting to investigate nonlinear robust formulations. Nonlinear SVM models are based on the use of kernel functions:

$$\begin{aligned}
\min_{\alpha, b, \xi} \quad & \frac{\lambda}{2} \alpha^\top K \alpha + \sum_{i=1}^L \xi_i \\
\text{subject to} \quad & y_i \left(\sum_{j=1}^L \alpha_j k(x_i, x_j) + b \right) + \xi_i \geq 1, \quad \forall i = 1, \dots, L, \\
& \xi_i \geq 0, \quad \forall i = 1, \dots, L,
\end{aligned} \tag{8}$$

where K is a kernel matrix whose elements are defined as

$$k_{ij} = \langle \phi(x_i), \phi(x_j) \rangle = k(x_i, x_j)$$

and k is a kernel function such as a polynomial, a Gaussian Radial Basis Function (RBF), or other specific choices of kernel functions [15].

One difficulty we are facing when designing robust counterparts formulations in the RKHS space is the understanding of how the uncertainty set is transformed through the implicit mapping chosen via the use of a specific kernel. The kernel function and its associated matrix K are known but its underlying mapping ϕ is unknown. It is therefore difficult to seek bounds of the image of the uncertainty set in the original space through ϕ . To address this issue, we investigate two methods : the Random Fourier Features (RFF) and the Nyström method. In both methods, the idea is to approximate the feature mapping ϕ by an explicit function $\tilde{\phi}$ from \mathbb{R}^n to \mathbb{R}^D ($n \leq D \leq L$) that will allow bounding of the image of uncertainties in the feature space. Specifically, if we can show that for all i in $\{1, \dots, L\}$ we can bound the uncertainty $\Delta\phi_i = \phi(x_i + \Delta x_i) - \phi(x_i)$ as follows:

$$\|R_i \Delta\phi_i\|_{\bar{p}} \leq \Gamma_i, \quad (9)$$

for some $L_{\bar{p}}$ norm in \mathbb{R}^D and where Γ_i is a constant, R_i is some matrix, both depending on the choice of feature map approximation, we are able to formulate the nonlinear robust SVM problem as follows

$$\min_{\zeta \in \mathbb{R}^D, b \in \mathbb{R}} \frac{\lambda}{2} \|\zeta\|^2 + \sum_{i=1}^L \max \left\{ 0, 1 - y_i (\langle \zeta, \tilde{\phi}(x_i) \rangle + b) + \Gamma_i \|R_i^\top \zeta\|_{\bar{q}} \right\}, \quad (10)$$

where \bar{q} is such that $\frac{1}{\bar{p}} + \frac{1}{\bar{q}} = 1$. The construction of this problem is similar to the construction of problem (4) in the linear case considering separation of approximate features $\tilde{\phi}(x_i)$ instead of data points x_i and the use of a robust loss function in the feature space as in (6).

As far as we know, the model (10) is the first model that is specifically designed to perform nonlinear classification on large and uncertain datasets. The decomposable structure of the objective allows, as for the linear case, the use of stochastic approximation methods such as the stochastic sub-gradient or stochastic proximal methods.

In the next 2 sections, we concentrate on deriving upper bounds of the form of (9) by the RFF and the Nyström methods.

4.1 Bounding feature uncertainties via Random Fourier Features

We propose first to make use of the so-called Random Fourier Features (RFF) to approximate the inner product $\langle \phi(x_i), \phi(x_j) \rangle$ by $\tilde{\phi}(x_i)^\top \tilde{\phi}(x_j)$ in a low dimensional Euclidian space \mathbb{R}^D , meaning that:

$$k(x, y) = \langle \phi(x), \phi(y) \rangle \approx \tilde{\phi}(x)^\top \tilde{\phi}(y), \quad \forall x, y \in \mathbb{R}^n.$$

The explicit knowledge of the randomized feature map $\tilde{\phi}$ helps in bounding the uncertainties in the randomized feature space. Additionally, as we expect D to be fairly low in practice, the technique should be able to handle large datasets as opposed to exact kernel methods that require the storage of a large and dense kernel matrix.

Mathematically, the technique relies on the Bochner theorem [13] which states that any shift invariant kernel $k(x, y) = k(x - y)$ is the inverse Fourier transform of a proper probability distribution P . If we define a complex-valued mapping $\hat{\phi} : \mathbb{R}^n \rightarrow \mathbb{C}$ by $\hat{\phi}_\omega(x) = e^{j\omega^\top x}$ and if we draw ω from P , the following holds:

$$k(x - y) = \int_{\mathbb{R}^n} P(\omega) e^{j\omega^\top (x-y)} d\omega = E \left[\hat{\phi}_\omega(x) \hat{\phi}_\omega(y)^* \right]$$

where $*$ denotes the complex conjugate. Furthermore, if instead we define

$$\tilde{\phi}_{\omega, \nu}(x) = \sqrt{2} \left[\cos(\omega^\top x + \nu) \right],$$

where ω is drawn from P and ν from the uniform distribution in $[0, 2\pi]$, we obtain an approximate real-valued random Fourier feature mapping that satisfies

$$E \left[\tilde{\phi}_{\omega, \nu}(x) \tilde{\phi}_{\omega, \nu}(y) \right] = k(x, y).$$

The method for constructing RFF to approximate a Gaussian kernel could therefore be summarized as follows:

1. Compute the Fourier transform P of the kernel k by taking:

$$P(\omega) = \frac{1}{2\pi} \int e^{-j\omega^\top \delta} k(\delta) d\delta;$$

2. Draw D i.i.d. samples $\omega_1, \dots, \omega_D$ in \mathbb{R}^n from P ;
3. Draw D i.i.d. samples ν_1, \dots, ν_D in \mathbb{R} from the uniform distribution on $[0, 2\pi]$;
4. The RFF $\tilde{\phi}_{\omega, \nu}(x)$ is given by


$$\tilde{\phi}_{\omega, \nu}(x) = \sqrt{\frac{2}{D}} \left[\cos(\omega_1^\top x + \nu_1), \dots, \cos(\omega_D^\top x + \nu_D) \right]^\top.$$

In [18], it has been shown that, when k is Gaussian, the following choice of features leads to lower variance and should be preferred in practice:

$$\tilde{\phi}_\omega(x) = \sqrt{\frac{2}{D}} \left[\cos(\omega_1^\top x), \sin(\omega_1^\top x), \dots, \cos(\omega_{D/2}^\top x), \sin(\omega_{D/2}^\top x) \right]^\top.$$

In Figure 1, a toy example is given. On the left side, one can see 4 nominal data points surrounded by their noisy observations. The 2 classes of points cannot be separated linearly. We construct random Fourier features with $D = 2$ on the

Figure 1: Random Fourier projections of uncertainties in a low dimensional feature space



right side. One can observe that the noisy observations are now projected on the circle and linear separation is possible.

We are now interested in the analysis of the image of uncertainties when mapped in \mathbb{R}^D through $\tilde{\phi}_\omega$. With the explicit knowledge of $\tilde{\phi}_\omega$ and if we define $\Delta\phi_{\omega,i} = \tilde{\phi}_\omega(x_i + \Delta x_i) - \tilde{\phi}_\omega(x_i)$ as the induced uncertainty around each random feature in the randomized feature space \mathbb{R}^D , we establish the following result:

Proposition 4.1 *For all i in $\{1, \dots, L\}$, let R_i be the block-diagonal matrix where the $D/2$ blocks of R_i are rotation matrices of angle $-\omega_j^\top x_i$ for j in $\{1, \dots, D/2\}$. We have*

$$\|R_i \Delta\phi_{\omega,i}\|_{\bar{p}} \leq \Gamma_i,$$

where Γ_i are constants depending on γ_i , $\Sigma_i^{1/2}$ and the $D/2$ vectors (ω_i) .

The proof of this result is as follows:

Proof: Let $c_{ij} = \cos(\omega_j^\top x_i)$ and $s_{ij} = \sin(\omega_j^\top x_i)$. Using trigonometric relations, the j -th block of $\Delta\phi_{\omega,i}$ can be written

$$\begin{aligned} (\Delta\phi_{\omega,i})_j &= \sqrt{\frac{2}{D}} \left(\cos(\omega_j^\top (x_i + \Delta x_i)) - \cos(\omega_j^\top x_i), \right. \\ &\quad \left. \sin(\omega_j^\top (x_i + \Delta x_i)) - \sin(\omega_j^\top x_i) \right)^\top \\ &= \sqrt{\frac{2}{D}} \left(c_{ij} \cos(\omega_j^\top \Delta x_i) - s_{ij} \sin(\omega_j^\top \Delta x_i) - c_{ij}, \right. \\ &\quad \left. s_{ij} \cos(\omega_j^\top \Delta x_i) + c_{ij} \sin(\omega_j^\top \Delta x_i) - s_{ij}, \right)^\top \\ &= \sqrt{\frac{2}{D}} P_{ij} \begin{pmatrix} \cos(\omega_j^\top \Delta x_i) - 1 \\ \sin(\omega_j^\top \Delta x_i) \end{pmatrix} \end{aligned}$$

where

$$P_{ij} = \begin{pmatrix} c_{ij} & -s_{ij} \\ s_{ij} & c_{ij} \end{pmatrix}$$

is the rotation matrix of angle $\omega_j^\top x_i$. Consequently, we obtain that

$$P_{ij}^\top (\Delta\phi_{\omega,i})_j = \sqrt{\frac{2}{D}} \begin{pmatrix} \cos(\omega_j^\top \Delta x_i) - 1 \\ \sin(\omega_j^\top \Delta x_i) \end{pmatrix}.$$

Using the following bounding schemes

$$|\cos(\theta) - 1| \leq \frac{\theta^2}{2}, \quad |\sin(\theta)| \leq |\theta|,$$

we have

$$\begin{aligned} |\cos(\omega_j^\top \Delta x_i) - 1| &\leq \frac{(\omega_j^\top \Delta x_i)^2}{2} \\ |\sin(\omega_j^\top \Delta x_i)| &\leq |\omega_j^\top \Delta x_i|. \end{aligned} \tag{11}$$

Using Hölder inequality, we can write

$$\begin{aligned} |\omega_j^\top \Delta x_i| &= |\omega_j^\top \Sigma_i^{1/2} \Sigma_i^{-1/2} \Delta x_i| \\ &\leq \left\| \omega_j^\top \Sigma_i^{1/2} \right\|_q \left\| \Sigma_i^{-1/2} \Delta x_i \right\|_p \\ &\leq \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \gamma_i, \end{aligned} \tag{12}$$

and

$$\begin{aligned} |\cos(\omega_j^\top \Delta x_i) - 1| &\leq \min \left(2, \frac{\left(\gamma_i \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \right)^2}{2} \right) \\ |\sin(\omega_j^\top \Delta x_i)| &\leq \min \left(1, \gamma_i \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \right). \end{aligned}$$

Introducing the notations

$$\begin{aligned} R_i &= \text{diag}(P_{i,1}^\top, \dots, P_{i,D/2}^\top), \\ \alpha_{i,j} &:= \min \left(2, \frac{\left(\gamma_i \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \right)^2}{2} \right) \end{aligned}$$

and

$$\beta_{i,j} := \min \left(1, \gamma_i \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \right),$$

we finally obtain the expected bound

$$\|R_i \Delta \Phi_{\omega, i}\|_{\bar{p}} \leq \Gamma_i$$

where

$$\Gamma_i := \begin{cases} \sqrt{\frac{2}{D}} \sum_{j=1}^{D/2} (\alpha_{i,j} + \beta_{i,j}) & \text{if } \bar{p} = 1; \\ \sqrt{\frac{4}{D}} \sum_{j=1}^{D/2} \alpha_{i,j} & \text{if } \bar{p} = 2; \\ \sqrt{\frac{2}{D}} \max \left(\max_{j=1, \dots, D/2} \alpha_{i,j}, \max_{j=1, \dots, D/2} \beta_{i,j} \right) & \text{if } \bar{p} = \infty. \end{cases}$$

□

4.2 Bounding feature uncertainties via the Nyström method

We now investigate another method. Note that in the following, we will use the notation $\|\cdot\|$ for the euclidean norm $\|\cdot\|_2$. When working with large datasets, the computation and storage of the entire kernel matrix K whose size is L^2 becomes difficult. For this reason, the Nyström method has been proposed to compute a low rank approximation of the kernel matrix [9, 21, 10, 11, 5, 4]. We can summarize the method as follows:

1. Choose an integer m between the number of attributes n and the size L of the dataset.
2. Choose m vectors $(\hat{x}_i)_{i=1, \dots, m}$ among the L training vectors.
3. Build the associated kernel matrix $\hat{K} \in \mathbb{R}^{m \times m}$ defined by

$$\hat{K}_{ij} := k(\hat{x}_i, \hat{x}_j).$$

4. Perform an eigenvalue decomposition of this matrix \hat{K} to obtain
 - its rank r ,
 - a $m \times r$ matrix $U_\sigma^{(r)} := (u_{\sigma,1}, \dots, u_{\sigma,r})$ such that $(U_\sigma^{(r)})^\top U_\sigma^{(r)}$ is the identity matrix,
 - and a diagonal matrix $\Lambda_\sigma^{(r)} := \text{diag}(\lambda_{\sigma,1}, \dots, \lambda_{\sigma,r})$

such that $\hat{K} U_\sigma^{(r)} = U_\sigma^{(r)} \Lambda_\sigma^{(r)}$.

5. The Nyström $\tilde{\phi}_\sigma(x)$ can be chosen as follows :

$$\tilde{\phi}_\sigma(x) = \left(\Lambda_\sigma^{(r)} \right)^{-1/2} \left(U_\sigma^{(r)} \right)^\top \hat{k}_\sigma(x),$$

with

$$\hat{k}_\sigma(x) = (k_\sigma(x, \hat{x}_1), k_\sigma(x, \hat{x}_2), \dots, k_\sigma(x, \hat{x}_m))^\top.$$

As before, we are interested by the analysis of the uncertainties when they are mapped to the approximate feature space induced by the Nyström method through $\tilde{\phi}_\sigma$. We define $\Delta\phi_{\sigma,i} = \tilde{\phi}_\sigma(x_i + \Delta x_i) - \tilde{\phi}_\sigma(x_i)$ as the uncertainty in the feature space \mathbb{R}^r and show the following result:

Proposition 4.2 For all i in $\{1, \dots, L\}$, let $R_i = \left(\Lambda_\sigma^{(r)}\right)^{1/2}$, we have

$$\|R_i \Delta\phi_{\sigma,i}\|_2 \leq \Gamma_i,$$

where Γ_i are constants depending on γ_i , $\Sigma_i^{1/2}$, σ and the m vectors $(\hat{x}_i)_{i=1, \dots, m}$.

Proof: Using the expression of the feature mapping constructed via the Nyström method, we have

$$\begin{aligned} \left(\Lambda_\sigma^{(r)}\right)^{1/2} \left(\tilde{\phi}_\sigma(x_i + \Delta x_i) - \tilde{\phi}_\sigma(x_i)\right) &= \left(U_\sigma^{(r)}\right)^\top \left(\hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i)\right) \\ &= \begin{pmatrix} u_{\sigma,1}^\top \left(\hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i)\right) \\ \vdots \\ u_{\sigma,r}^\top \left(\hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i)\right) \end{pmatrix}. \end{aligned}$$

Consequently, we can write

$$\begin{aligned} \left\| \left(\Lambda_\sigma^{(r)}\right)^{1/2} \left(\tilde{\phi}_\sigma(x_i + \Delta x_i) - \tilde{\phi}_\sigma(x_i)\right) \right\|_2^2 &= \sum_{j=1}^r \left(u_{\sigma,j}^\top \left(\hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i)\right) \right)^2 \\ &\leq \sum_{j=1}^r \left\| \hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) \right\|_2^2 \\ &= r \left\| \hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) \right\|_2^2. \end{aligned}$$

We are now interested in computing an upper bound of the left hand side of the above expression. Since

$$\hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) = \begin{pmatrix} k_\sigma(x_i + \Delta x_i, \hat{x}_1) - k_\sigma(x_i, \hat{x}_1) \\ \vdots \\ k_\sigma(x_i + \Delta x_i, \hat{x}_m) - k_\sigma(x_i, \hat{x}_m) \end{pmatrix}$$

we get

$$\left\| \hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) \right\|_2^2 = \sum_{j=1}^m \left(k_\sigma(x_i + \Delta x_i, \hat{x}_j) - k_\sigma(x_i, \hat{x}_j) \right)^2.$$

If the uncertainties are defined with the choice of Euclidian norm (meaning $p = q = 2$), if we also define the kernel k_σ with the Euclidian norm (meaning

$k_\sigma(x, z) := e^{-\frac{\|x - z\|_2^2}{2\sigma^2}}$), and observing that

$$\begin{aligned}
k_\sigma(x_i + \Delta x_i, \hat{x}_j) - k_\sigma(x_i, \hat{x}_j) &= e^{-\frac{\|x_i + \Delta x_i - \hat{x}_j\|^2}{2\sigma^2}} - e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \\
&= e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \left(e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j) + \|\Delta x_i\|^2}{2\sigma^2}} - 1 \right),
\end{aligned}$$

we can obtain

$$\begin{aligned}
(k_\sigma(x_i + \Delta x_i, \hat{x}_j) - k_\sigma(x_i, \hat{x}_j))^2 &= \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \left(e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j) + \|\Delta x_i\|^2}{2\sigma^2}} \right)^2 \\
&\quad - 2 \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j) + \|\Delta x_i\|^2}{2\sigma^2}} \\
&\quad + \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \\
&= \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \left(e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j)}{2\sigma^2}} \right)^2 \left(e^{-\frac{\|\Delta x_i\|^2}{2\sigma^2}} \right)^2 \\
&\quad - 2 \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j)}{2\sigma^2}} e^{-\frac{\|\Delta x_i\|^2}{2\sigma^2}} \\
&\quad + \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2.
\end{aligned}$$

We have easily $\left(e^{-\frac{\|\Delta x_i\|^2}{2\sigma^2}} \right)^2 \leq 1$. We are now aiming at bounding from below

the terms $e^{-\frac{\|\Delta x_i\|^2}{2\sigma^2}}$ and $e^{-\frac{2\Delta x_i^\top (x_i - \hat{x}_j)}{2\sigma^2}}$. Using $\|\Sigma_i^{-1/2} \Delta x_i\|_2 \leq \gamma_i$, we can write

$$0 \leq \|\Delta x_i\| = \|\Sigma_i^{1/2} \Sigma_i^{-1/2} \Delta x_i\| \leq \|\Sigma_i^{1/2}\| \|\Sigma_i^{-1/2} \Delta x_i\| \leq \gamma_i \|\Sigma_i^{1/2}\|.$$

which consequently leads to, $-\gamma_i^2 \|\Sigma_i^{1/2}\|^2 \leq -\|\Delta x_i\|^2 \leq 0$ and

$$e^{-\frac{\gamma_i^2 \|\Sigma_i^{1/2}\|^2}{2\sigma^2}} \leq e^{-\frac{\|\Delta x_i\|^2}{2\sigma^2}} \leq 1.$$

Now, writing $\Delta x_i^\top(x_i - \hat{x}_j) = (x_i - \hat{x}_j)^\top \Sigma_i^{1/2} \Sigma_i^{-1/2} \Delta x_i$, we have $|\Delta x_i^\top(x_i - \hat{x}_j)| \leq \gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|$ and thus

$$e^{-\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2}} \leq e^{-\frac{2\Delta x_i^\top(x_i - \hat{x}_j)}{2\sigma^2}} \leq e^{\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2}}.$$

Using these bounds, we can write

$$\begin{aligned} (k_\sigma(x_i + \Delta x_i, \hat{x}_j) - k_\sigma(x_i, \hat{x}_j))^2 &\leq \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \left(\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2} \right)^2 \\ &\quad - 2 \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 e^{-\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2}} e^{-\frac{\gamma_i^2 \|\Sigma_i^{1/2}\|^2}{2\sigma^2}} \\ &\quad + \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \end{aligned}$$

and come back to

$$\begin{aligned} \left\| \hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) \right\|_2^2 &\leq \sum_{j=1}^m \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 \left(1 + \left(\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2} \right)^2 \right) \\ &\quad - 2e^{-\frac{\gamma_i^2 \|\Sigma_i^{1/2}\|^2}{2\sigma^2}} \sum_{j=1}^m \left(e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}} \right)^2 e^{-\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2}} \end{aligned}$$

or

$$\left\| \hat{k}_\sigma(x_i + \Delta x_i) - \hat{k}_\sigma(x_i) \right\|_2^2 \leq \sum_{j=1}^m k_{ij}^2 \left(\frac{1}{\tau_{ij}^2} + 1 \right) - 2\rho_i \sum_{j=1}^m k_{ij}^2 \tau_{ij}$$

where we have denoted

$$k_{ij} = e^{-\frac{\|x_i - \hat{x}_j\|^2}{2\sigma^2}}, \quad \tau_{ij} = e^{-\frac{2\gamma_i \|\Sigma_i^{\top/2}(x_i - \hat{x}_j)\|}{2\sigma^2}}, \quad \rho_i = e^{-\frac{\gamma_i^2 \|\Sigma_i^{1/2}\|^2}{2\sigma^2}}.$$

Finally, we have obtained

$$\left\| \left(\Lambda_\sigma^{(r)} \right)^{1/2} \left(\tilde{\phi}_\sigma(x_i + \Delta x_i) - \tilde{\phi}_\sigma(x_i) \right) \right\|_2^2 \leq \Gamma_i^2$$

with

$$\Gamma_i = \sqrt{r \left(\sum_{j=1}^m k_{ij}^2 \left(\frac{1}{\tau_{ij}^2} + 1 \right) - 2\rho_i \sum_{j=1}^m k_{ij}^2 \tau_{ij} \right)}.$$

□

5 A word on the tightness of bounding schemes

We would like here to analyze under which conditions the bounds we have proposed are tight. Ensuring some tightness will avoid over estimating the uncertainties in the approximate feature space that could make class separation more difficult.

Consider the bounding scheme from Section 4.1 and where k is chosen as the Gaussian kernel, meaning that $k_\sigma(x, z) := e^{-\frac{\|x - z\|^2}{2\sigma^2}}$. The method relies on the inequalities (11) that ensure that

$$\begin{aligned} \left| \cos(\omega_j^\top \Delta x_i) - 1 \right| &\leq \frac{(\omega_j^\top \Delta x_i)^2}{2} \\ \left| \sin(\omega_j^\top \Delta x_i) \right| &\leq \omega_j^\top \Delta x_i. \end{aligned}$$

These bounds are rather tight if the angle $|\omega_j^\top \Delta x_i|$ is smaller than some given θ_{max} . To ensure (with high probability) that the angle will remain below θ_{max} we state and proof the following result:

Proposition 5.1 *Assume that $\Sigma_i^{1/2} = \Sigma^{1/2}$ where $\Sigma^{1/2}$ is a constant diagonal matrix. When $q = 2$, if*

$$\sigma \geq \frac{3\gamma_i \|\Sigma^{1/2}\|_F}{\theta_{max}},$$

with high probability (greater than 0.997), we have $|\omega_j^\top \Delta x_i| \leq \theta_{max}$.

Proof: Remember that ω_j are distributed according to the Fourier transform of the kernel k , therefore we have $\omega_j \sim \mathcal{N}(0, \frac{1}{\sigma^2})$ and we know that $\omega_j \in [-\frac{3}{\sigma}, \frac{3}{\sigma}]$ with high probability (greater than 0.997). Additionally, we have shown in (12) that for all j in $\{1, \dots, D/2\}$,

$$|\omega_j^\top \Delta x_i| \leq \left\| \Sigma_i^{\top/2} \omega_j \right\|_q \gamma_i. \quad (13)$$

When $q = 2$ and $\Sigma_i^{1/2} = \Sigma^{1/2} = \text{diag}(s)$, we have

$$\left\| \Sigma_i^{\top/2} \omega_j \right\|_2^2 = \sum_{k=1}^n (s_k \omega_{jk})^2 \leq \frac{9 \|\Sigma^{1/2}\|_F^2}{\sigma^2},$$

and therefore, when we take

$$\frac{9\gamma_i^2\|\Sigma^{1/2}\|_F^2}{\sigma^2} \leq \theta_{max}^2, \tag{14}$$

using (13) we have, with high probability, that $|\omega_j^\top \Delta x_i| \leq \theta_{max}$. The inequality (14) can also be written as $\sigma \geq \frac{3\gamma_i\|\Sigma^{1/2}\|_F}{\theta_{max}}$. \square

The result in Proposition 5.1 actually says that there is a trade-off between bounding the uncertainties in the approximate feature space and separating correctly the uncertain features. For the latter, we would like to achieve relatively small (but not too small) values of σ . On the other hand to ensure tightness of our bounding scheme, the RFF method requires that σ should be taken sufficiently large. Therefore, the right trade-off for choosing the σ value is depending on data. The condition in Proposition 5.1 should then be incorporated in the definition of the σ -grid that one usually defines when designing a cross-validation procedure in practice.

References

- [1] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Mathematical Programming*, 95(1):3–51, 2003.
- [2] A. Ben-Tal, L. El Ghaoui, and A.S. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, October 2009.
- [3] Aharon Ben-Tal, Sahely Bhadra, Chiranjib Bhattacharyya, and J. Saketha Nath. Chance constrained uncertain classification via robust optimization. *Mathematical Programming*, 127(1):145–173, 2011.
- [4] Lo-Bin Chang, Zhidong Bai, Su-Yun Huang, and Chii-Ruey Hwang. Asymptotic error bounds for kernel-based Nyström low-rank approximation matrices. *J. Multivariate Anal.*, 120:102–119, 2013.
- [5] Anna Choromanska, Tony Jebara, Hyungtae Kim, Mahesh Mohan, and Claire Monteleoni. Fast spectral clustering via the Nyström method. In *Algorithmic learning theory*, volume 8139 of *Lecture Notes in Comput. Sci.*, pages 367–381. Springer, Heidelberg, 2013.
- [6] Nicolas Couellan and Wenjuan Wang. On the convergence of stochastic bi-level gradient methods. *Eprint - Optimization Online*, 2016.
- [7] Nicolas Couellan and Wenjuan Wang. Uncertainty-safe large scale support vector machines. *Comput. Statist. Data Anal.*, 109:215–230, 2017.
- [8] Nello Cristianini and John Shawe-Taylor. *An introduction to support vector machines and other kernel-based learning methods. Repr.* Cambridge: Cambridge University Press, repr. edition, 2001.

- [9] Alex Gittens and Michael W. Mahoney. Revisiting the Nyström method for improved large-scale machine learning. *J. Mach. Learn. Res.*, 17:1–65, 2016.
- [10] Darren Hornrighausen and Daniel J. McDonald. On the Nyström and column-sampling methods for the approximate principal components analysis of large datasets. *J. Comput. Graph. Statist.*, 25(2):344–362, 2016.
- [11] Mu Li, Wei Bi, James T. Kwok, and Bao-Liang Lu. Large-scale Nyström kernel matrix approximation using randomized SVD. *IEEE Trans. Neural Netw. Learn. Syst.*, 26(1):152–164, 2015.
- [12] MOSEK-ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 7.1 (Revision 28).*, 2015.
- [13] Ali Rahimi and Ben Recht. Random features for large-scale kernel machines. In *Neural Information Processing Systems*, 2007.
- [14] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [15] John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, New York, NY, USA, 2004.
- [16] P.K. Shivaswamy, C. Bhattacharyya, and A. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7:1283–1314, 2006.
- [17] J.F. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 1112:625653, 1999.
- [18] Dougal J Sutherland and Jeff Schneider. On the error of random fourier features. *arXiv preprint arXiv:1506.02785*, 2015.
- [19] T. Trafalis and R. Gilbert. Robust support vector machines for classification and computational issues. *Optimization Methods and Software*, 22(1):187–198, 2007.
- [20] T. B. Trafalis and R. C. Gilbert. Robust classification and regression using support vector machines. *European Journal of Operational Research*, 173(3):893–909, 2006.
- [21] Aleksandar Trokicić. Approximate spectral learning using Nyström method. *Facta Univ. Ser. Math. Inform.*, 31(2):569–578, 2016.
- [22] Vladimir Naoumovitch Vapnik. *Statistical learning theory*. Adaptive and learning systems for signal processing, communications, and control. Wiley, New York, 1998.

- [23] Silvia Villa, Lorenzo Rosasco, and Bang Cong Vu. Learning with stochastic proximal gradient. In *Workshop on Optimization for Machine Learning. At the conference on Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [24] P. Xanthopoulos, P. Pardalos, and T.B. Trafalis. *Robust Data Mining*. SpringerBriefs in Optimization. Springer, 2012.