One-Class Classification Using ℓ_p -Norm Multiple Kernel Fisher Null-Space

Shervin Rahimzadeh Arashloo

Abstract—The paper addresses the one-class classification (OCC) problem and advocates a one-class multiple kernel learning (MKL) approach for this purpose. To this aim, based on the Fisher null-space one-class classification principle, we present a multiple kernel learning algorithm where an ℓ_p -norm constraint ($p \geq 1$) on kernel weights is considered. We cast the proposed one-class MKL task as a min-max saddle point Lagrangian optimisation problem and propose an efficient method to solve it. An extension of the proposed one-class MKL approach is also considered where several related one-class MKL tasks are learned concurrently by constraining them to share common kernel weights.

An extensive assessment of the proposed method on a range of data sets from different application domains confirms its merits against the baseline and several other algorithms.

Index Terms—One-class classification, multiple kernel learning, one-class Fisher null transformation, ℓ_p -norm regularisation.

I. INTRODUCTION

ONE-class classification (OCC) corresponds to recognition of patterns that meet a particular condition identified as normal, and discerning them from any irregular observation diverging from normality, known as anomalies, abnormalities, novelties, etc. [1], [2]. Unlike the common multi-class classification paradigm, OCC mainly rests on observations from a single (typically normal/target) class for model construction. A one-class design scheme may be favoured in practice due to different reasons such as imbalanced data [3] or difficulty in accessing representative training samples from the non-target class. The limitations in obtaining prototypical training observations may arise, for instance, because of the high cost associated with collecting samples from the non-target class, the unpredictable form of test observations, the openness of a recognition problem, etc. Such situations may arise in a variety of applications including biometric presentation attack detection [4], health care [5], surveillance [6], intrusion detection [7], safety-critical systems [8], fraud detection [9], insurance [10], etc. As such, OCC serves as an indispensable ingredient to a diverse range of real-world systems.

There exist different studies on OCC, resulting in a plethora of different methods, including the density methods, the reconstruction-based techniques, and the boundary approaches [11]. A popular kernel-based one-class classification method is the support vector data description (SVDD) approach [12], [13] which draws on the SVM classifier and estimates a spherical boundary to enclose positive training samples which

is then used to classify test observations. A closely related method to that of SVDD is the one-class SVM (OCSVM) approach [14] where the positive instances are split from the origin using a maximum margin hyperplane. Some other instances of successful kernel-based OCC methods are the GP (Gaussian Process) method [15] and the one-class KPCA (Kernel Principal Component Analysis) approach [16]. Among others, the work in [17], [18] present a unique approach called the *one-class Fisher null-space* method. In this method, the one-class classifier operates on an adaptation of the Fisher classification principle to the one-class setting, yielding a theoretically optimal Fisher criterion value. While conventionally the Fisher classifier requires both positive and negative samples for training, in the one-class Fisher null-space approach, the requirement for the existence of negative training observations is circumvented by presuming the origin as an artificial negative sample and training the classifier using only positive/target observations. Through a formulation in a RKHS (reproducing kernel Hilbert space) and by using non-linear kernel functions, the one-class Fisher null-space method exhibits flexibility to deal with inherently non-linear data. The promising performance of this approach as compared with several other alternatives in different applications [19], [18], [17], [20] has led to further developments and improvements of this method [21].

The formulation of the one-class Fisher null-space approach in the kernel space enables decoupling of learning the classifier from data representations while opening the possibility to benefit from powerful machine learning techniques. Nevertheless, similar to other kernel-based methods, selecting an appropriate kernel function is an important design consideration in the Fisher null-space OCC method as the kernel characterises the embedding of the observations in the feature space, which directly impacts on the performance. While, preferably, the embedding should be learnt directly from the data, in kernel-based methods, this challenging problem is usually relaxed as searching for an optimal combination of multiple kernels, each capturing a distinct view of the problem. Accessibility of multiple kernel candidates for a specific problem may arise, for example, due to multiple data representations, heterogeneous data sources, or as a result of using different kernel functions. The availability of multiple kernels and the desire to benefit from multiple information sources for an improved performance motivates the studies concerning an automatic combination of several kernels, referred to as multiple kernel learning (MKL) [22], [23], [24], [25], [26], [27], [28]. Very often, the composite kernel is expressed in terms of a linear combination of multiple base kernels, rendering the MKL problem as one of searching for the optimal combination

S.R. Arashloo is with the Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, Guildford, Surrey, GU2 7XH, U.K. e-mail: s.rahimzadeh@surrey.ac.uk.

Manuscript received April ?, ?; revised ?, ?.

weights under the constraint that a valid combined kernel is derived.

Among others, the choice of a particular regularisation constraint on kernel weights is an important design issue that directly affects the performance of an MKL algorithm. Historically, an ℓ_1 -norm regularisation that leads to sparse coefficients has been one of the popular choices [29]. Nevertheless, sparse MKL models have been frequently observed to be outperformed by their non-sparse counterparts as they tend to discard some useful representations from the model [30], [31], [32]. In multi-class scenarios and in response to the quest for a model with the potential of improving the accuracy over a uniform-weight kernel combination rule, an ℓ_p - or a mixed-norm regularisation may be considered [30], [31], [33]. While in the multi-class case alternative MKL formulations based on the SVM or the Fisher discriminant analysis exist [22], [30], [31], [34], for OCC, the MKL methods appear to be relatively scarce and formulated in terms of one-class SVM's with no earlier study addressing the one-class MKL problem within the *Fisher* framework. Moreover, in the multi-class classification paradigm, different regularisation constraints including the ℓ_p - and the mixed-norm have been considered whereas in the one-class setting, the focus has been primarily on an ℓ_1 -norm regularisation [35], [36], [37].

In this work, we consider the one-class MKL problem and address it in the context of the kernel Fisher null-space method [17], [18], and in particular, in terms of its regression-based reformulation [21] for one-class classification. We formulate the one-class MKL problem by considering a general ℓ_p -norm constraint ($p \geq 1$) on kernel weights, providing the model with the flexibility to cope with problems with different sparsity characteristics. More specifically, the inclusion of a regularisation parameter in the proposed approach allows the method to tune into the intrinsic sparsity of a specific one-class task to improve the performance. Additionally, we illustrate that when several one-class problems are presumed to be related in terms of their kernel space representations, the proposed approach may be naturally extended to *jointly* learn multiple OCC MKL problems by constraining them to share common kernel weights.

A. Contributions

The major contributions of the present work are as follows.

- Based on the Fisher null-space one-class classification method [17], [18], [21], we propose an ℓ_p -norm one-class multiple kernel learning approach (" ℓ_p MK-FN") for one-class classification and formulate it as a saddle point Lagrangian optimisation problem.
- We introduce an effective method to solve the min-max optimisation problem associated with the ℓ_p MK-FN method.
- We extend the ℓ_p MK-FN method to learn several related one-class problems jointly (" ℓ_p Joint MK-FN") by coupling different OCC problems in terms of kernel weights.
- We conduct a thorough evaluation of the proposed approach on a range of different databases from different application domains and provide a comparison to the baseline and several other one-class multiple kernel fusion

methods, including the SVM-based MKL techniques and end-to-end learning one-class deep methods to illustrate the merits of the proposed method.

B. Organisation

The remainder of the article is structured as detailed next. In Section II, a review of the relevant work on one-class multiple kernel learning is provided. In Section III, we first provide a brief background on the kernel Fisher null-space approach [17], [18] for one-class classification followed by a description of its regression-based reformulation for OCC as considered in [21]. In Section IV, we present our new ℓ_p -norm multiple kernel Fisher null-space (ℓ_p MK-FN) approach for one-class classification. In Section VI, an extension of the ℓ_p MK-FN method for *joint* multiple kernel learning of several related one-class problems is introduced. The results of an extensive experimental evaluation of the proposed method along with a comparison against other techniques on different datasets is presented in Section VII. Finally, in Section VIII, conclusions are drawn.

II. RELATED WORK

As the focus of this study is on one-class classification, we skip a review of the multi-class multiple kernel learning methods. One may consult [22], [30], [25] for a taxonomy and comparison of multi-class multiple kernel learning algorithms.

For one-class classification, there exist different MKL algorithms. As an instance, in [35], the MKL problem is addressed through a weighted 2-norm regularisation constraint on kernel weights that promotes sparse weights. The authors proposed an algorithm called SimpleMKL to solve the formulated MKL problem. Although the method is mainly developed for binary classification based on an SVM formulation, the extension to the one-class setting for an ℓ_1 -norm OCSVM [14] is also discussed. Other work [36] presents a modification of the SVDD approach to multiple kernels and applies the methodology presented in the SimpleMKL approach to enforce an ℓ_1 -norm regularisation on kernel weights. By observing that in an SVM formulation a larger number of support vectors lead to tighter solution boundaries, the study also considers the so-called *slim* variants of the multiple kernel SVDD and OCSVM methods by modifying the objective function in a way that tighter solutions are favoured over loose ones. For comparison purposes, an equivalent method is also developed for the One-Class SVM and compared for 3D shape analysis and detecting outliers where the proposed approaches are deployed for both supervised and unsupervised outlier detection. The work in [37] focuses on semi-infinite linear programming formulations of the ℓ_1 -norm SVM-based multi-class MKL problem. The proposed formulation is shown to be solved efficiently via recycling the standard support vector machine implementations. Possible generalisations of this approach to handle a wider set of problems including regression and one-class classification are discussed. Nevertheless, no experimental analysis for one-class classification is provided in this work. In a different study [38], and in contrary to the previous studies that assign a common weight to the entire

observation space, instead of assigning a common weight spanning the whole observation space, a localized multiple kernel learning algorithm, comprised of a learning method and a parametrised function for assigning local weights to the kernels is presented. The two ingredients of the algorithm are then trained using an alternating optimisation approach in a coupled fashion. An ℓ_1 -norm regularisation is implicitly assumed by choosing specific gating functions for assigning local weights. The work in [31] studies the MKL problem with a special focus on an SVM-based formulation and provide insights and connections between different MKL formulations and propose two different optimisation strategies for the MKL problem subject to an ℓ_p -norm regularisation constraint. Although the main focus of the work in [31] is on two-/multi-class classification, one-class classification is also discussed as a special case. Yet, no experimental analysis for the one-class scenario is carried out.

Despite that fact that in some of the studies mentioned above generic loss functions are discussed [31], they primarily concentrate on the hinge loss function. In this regard, the corresponding multiple kernel learning algorithms principally correspond to multiple kernel support vector machines. In comparison to support vector machines where the soft margin is maximised, in the Fisher discriminant analysis the ratio of the between- and the within-class scatter in a subspace is maximised. The outstanding classification performance of the Fisher classification principle as compared with other competitors, including the support vector machines has established the Fisher classifier a favourable choice for many classification problems over multiple decades. Its extension to the null-space variant which attains the theoretically optimal value of the Fisher ratio has made it an even more appealing classification approach not only for multi-class problems but also for one-class classification tasks by providing comparable or even superior performance compared with not only SVM but also deep end-to-end learning methods when operating on pre-trained deep representations.

Considering the body of existing work on MKL for OCC, first, one notices that the existing one-class classification MKL methods focus on SVM-based formulations. In contrast, in this study, a Fisher classification principle for one-class multiple kernel learning is considered. Moreover, we consider a general ℓ_p -norm constraint on kernel weights in contrast to the majority of other studies focusing on a fixed-norm (typically an ℓ_1 -norm) regularisation in the one-class setting. And last but not the least, we show that the proposed method can be naturally extended to jointly learn several related OCC problems and discuss its advantages. In this respect, there exists no earlier study on *joint* one-class multiple kernel learning for several related one-class problems. As will be discussed further in the subsequent sections, motivated by the success of the one-class Fisher null approach for OCC, the current study explores the ℓ_p -norm MKL paradigm for the one-class Fisher null classification principle and presents an MKL algorithm for this purpose. In this formalism, the current study illustrates that an optimal multiple kernel combination of pre-trained features obtained from the common CNNs currently in use, yields a very good performance for OCC while it is also applicable to

non-deep features.

III. PRELIMINARIES

In this section, first, an overview of the one-class Fisher null-space approach [17], [18] is presented. Then, we review how the Fisher null-space method is posed as a regression problem in our earlier study [21] for one-class classification.

A. One-Class Fisher Null-Space

The one-class kernel Fisher null-space approach for one-class classification [17], [18] is based on the Fisher classification principle. In a Fisher classifier, one seeks the maximiser of the objective function $G(v)$:

$$\arg \max_v G(v) = \arg \max_v \frac{v^\top S_b v}{v^\top S_w v} \quad (1)$$

S_w and S_b stand for the within-class and the between-class scatter matrices, respectively, while v represents one axis of the subspace. The sole maximiser of the problem in Eq. 1 is given as the leading eigenvector of the generalised eigenvalue problem

$$S_b v = \lambda S_w v \quad (2)$$

After deriving v , a sample x is projected onto the Fisher subspace as

$$y = v^\top x \quad (3)$$

The theoretically *optimal* discriminant that maximises the criterion function $G(v)$ is the one that leads to a positive between-class scatter and a zero within-class scatter, i.e.:

$$\begin{aligned} v^\top S_b v &> 0 \\ v^\top S_w v &= 0 \end{aligned} \quad (4)$$

The one-class null-space Fisher method [17], [18] is an adaptation of the Fisher criterion to the one-class setting that solves the Fisher criterion to optimality. In a one-class setting where only positive training instances are available, in the one-class Fisher null-space method, similar to some other studies [14], the non-target (negative) class is characterised by a single hypothetical example lying at the origin. The goal is then to find a mapping such that all target/normal training instances are projected onto a similar point distinct from the origin. Such a projection is formulated in a reproducing kernel Hilbert space (RKHS) using non-linear kernel functions to achieve the flexibility to handle data with inherently non-linear nature. The remarkable performance of this method for one-class classification as compared with several other alternatives is verified in different studies [21], [18], [17], [20], [19].

B. Regression-Based One-Class Fisher Null-Space

Finding the optimal projection corresponding to the one-class Fisher null-space method incorporates eigen-decomposition of dense matrices [17], [18]. As the corresponding eigen-decompositions are computationally demanding, a reformulation of this method in terms of regression in the RKHS is considered in [21]. As discussed in [21],

a reformulation of the Fisher null-space method in terms of regression for OCC avoids the computationally intensive eigen-decomposition of matrices and paves the way to impose different regularisations on the solution to improve the generalisation performance. A brief overview of the reformulation of the kernel Fisher null-space method in terms of regression for one-class classification, as discussed in [21], is provided next.

Let us assume that there exist n positive (target) training observations x_i 's. As noted earlier, the goal in the one-class kernel Fisher null-space approach is to map all the target training observations x_i 's onto the same point in the subspace, distinct from the origin. Without loss of generality, let us assume that all samples are to be mapped onto point 1. To realise such a projection, instead of following an eigen-decomposition approach as practised in [17], [18], the following proposition [21] may be used:

Proposition 1 *Let $\phi(x_i)$ denote the feature vector associated with x_i in the RKHS. If $\mathbf{w}$ represents the optimal answer to*

$$\min_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n (1 - \phi(\mathbf{x}_i)^\top \mathbf{w})^2 \quad (5)$$

then the projection $\phi(\cdot)^\top \mathbf{w}$ maps all target training samples onto point 1, i.e. $\phi(\mathbf{x}_i)^\top \mathbf{w} = 1, i = 1, \dots, n$.

A proof is provided in Section I of the supplementary material. Since through projection $\phi(\cdot)^\top \mathbf{w}$ all target observations are mapped onto point 1, the within-class variance for the positive class shall be zero. Moreover, the within-class variance of the negative training set is also zero since the negative class is represented using a single artificial observation lying at the origin. Hence, the total within-class scatter will be zero. Furthermore, as the projection of the positive training instances (i.e. 1) is distinct from that of the negative training set (i.e. the origin), the projection above provides a between-class scatter which is positive. Consequently, it corresponds to a Fisher null-space projection for one-class classification (further details are given in Section II of the supplementary material).

In [21], a regularisation scheme is employed to constrain the solution of Eq. 5. In essence, regularising the solution enforces some limitation on the functional space via incorporating a penalty term so that some regions of the solution space are less favoured. In particular, in [21] a Tikhonov regularisation is deployed to penalise coefficients with larger magnitude to derive a more parsimonious solution. The impact of a Tikhonov regularisation is to obtain a smoother projection function [39] with better generalisation capabilities [21]. A Tikhonov regularisation on the solution of Eq. 5 is applied as

$$\min_{\mathbf{w}} \|\mathbf{w}\|_2^2 + \frac{\theta}{n} \sum_{i=1}^n (1 - \phi(\mathbf{x}_i)^\top \mathbf{w})^2 \quad (6)$$

where θ is a regularisation parameter. The projection $\phi(\cdot)^\top \mathbf{w}$ where $\mathbf{w}$ minimises the problem above corresponds to the Tikhonov-regularised one-class kernel Fisher null-space projection for OCC and is shown to outperform several other one-class classification methods in different scenarios [21].

In kernel-based methods, it is common practice and typically more convenient to work in the dual space, as discussed in the next section.

IV. ℓ_p -NORM MULTIPLE KERNEL ONE-CLASS FISHER NULL-SPACE

The one-class multiple kernel learning algorithm proposed in this work is based on the dual form of the regularised kernel Fisher null-space method. For deriving the dual of the optimisation problem in Eq. 6, let $\delta = n/\theta$. The unconstrained optimisation problem in Eq. 6 may then be written as a constrained optimisation problem as

$$\begin{aligned} \min_{\mathbf{w}} \mathbf{w}^\top \mathbf{w} + \frac{1}{\delta} \|\zeta\|_2^2 \\ \text{s.t. } \zeta = \mathbf{1} - \Phi(\mathbf{X})^\top \mathbf{w} \end{aligned} \quad (7)$$

where $\Phi(\mathbf{X})$ stands for the matrix of training samples in the RKHS and $\mathbf{1}$ is an n -vector of ones. The Lagrangian for the constrained optimisation problem above may be formed as

$$\mathcal{L} = \mathbf{w}^\top \mathbf{w} + \frac{1}{\delta} \|\zeta\|_2^2 + \rho^\top (\mathbf{1} - \Phi(\mathbf{X})^\top \mathbf{w} - \zeta) \quad (8)$$

where ρ is the Lagrange multiplier. In order to find the optimal $\mathbf{w}$, one needs to solve the dual problem. For this purpose, one first needs to minimise $\mathcal{L}$ in $\mathbf{w}$ and ζ and then maximize it in ρ . In order to minimise $\mathcal{L}$ w.r.t. $\mathbf{w}$ and ζ , the corresponding partial derivatives are set to zero. Omitting the intermediate steps, the dual problem as

$$\max_{\alpha} -\alpha^\top \mathbf{K} \alpha - \delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} \quad (9)$$

where in the last equation the definition of a kernel matrix ($\mathbf{K}$) is used. In this work, the composite kernel is presumed to be a linear combination of non-negatively weighted kernels, subject to an ℓ_p -norm constraint on kernel weights. Assuming J base kernels to be combined, the kernel matrix $\mathbf{K}$ in Eq. 9 will be replaced by $\sum_{j=1}^J \beta_j \mathbf{K}_j$ where β_j 's denote the non-negative kernel weights. Replacing $\mathbf{K}$ by $\sum_{j=1}^J \beta_j \mathbf{K}_j$ and optimising over β (a vector collection of kernel weights) subject to the non-negativity and p -norm constraints, the optimisation problem for the proposed ℓ_p -norm multiple kernel one-class Fisher null-space method shall be derived as a saddle point problem:

$$\begin{aligned} \min_{\beta} \max_{\alpha} \quad & -\alpha^\top (\sum_{j=1}^J \beta_j \mathbf{K}_j) \alpha - \delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} \\ \text{s.t. } \quad & \beta \geq 0, \|\beta\|_p^p \leq 1 \end{aligned} \quad (10)$$

where $\beta \geq 0$ is meant element-wise. Under the condition that each individual kernel matrix $\mathbf{K}_j$ is positive semidefinite, the non-negativity constraint $\beta \geq 0$ guarantees that the combined kernel matrix is also positive semidefinite and is thus a valid kernel matrix while $\|\beta\|_p^p \leq 1$ imposes a general ℓ_p -norm constraint on kernel weights. As noted previously, the introduction of a p -norm regularisation on kernel weights is beneficial as it allows the model to adjust to the underlying sparsity structure of the problem under consideration. This is in contrast to a fixed-norm regularisation which may not be able to effectively represent the intrinsic sparsity of the

Fig. 1. The unit ℓ_p -norm balls (in 2D) for $p \geq 1$ in the first quadrant.

problem. Similar observations have been made in the multi-class scenario where a variable-norm regularisation has been found to be advantageous compared to fixed-norm solutions [31], [30].

V. OPTIMISATION

Let us assume $\mathbf{u}$ a vector with J elements whose j^{th} element is $u_j = \alpha^\top \mathbf{K}_j \alpha$. The saddle point problem in Eq. 10 may now be expressed as

$$\begin{aligned} \min_{\beta} \max_{\alpha} -\delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} - \beta^\top \mathbf{u} \\ \text{s.t. } \beta \geq 0, \|\beta\|_p^p \leq 1 \end{aligned} \quad (11)$$

For fixed α , the objective function is linear in β and the constraint set for $p \geq 1$ (the focus of this study) forms a closed convex set, Fig. 1. Consequently, for α being fixed, the cost function is convex in β . Moreover, when β is fixed, the function would be concave in α . As a result, according to the generalised minimax theorem [40], [41], Eq. 11 may be equivalently written as

$$\begin{aligned} \max_{\alpha} -\delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} + \min_{\beta} -\beta^\top \mathbf{u} \\ \text{s.t. } \beta \geq 0, \|\beta\|_p^p \leq 1 \end{aligned} \quad (12)$$

That is, one may exchange the order of optimisation w.r.t. β and α . In order to optimise the cost function in β , we form minimisation sub-problem's Lagrangian as

$$\mathcal{L} = -\beta^\top (\mathbf{u} + \boldsymbol{\mu}) + \gamma (\|\beta\|_p^p - 1) \quad (13)$$

where $\gamma \geq 0$ and $\boldsymbol{\mu} \geq 0$ are Lagrange multipliers. It may be easily verified that the Slater's condition holds and the KKT conditions are

$$\begin{aligned} \nabla_{\beta} \mathcal{L} &= -(\mathbf{u} + \boldsymbol{\mu}) + \gamma p |\beta|^{p-1} \odot \text{sign}(\beta) = 0 \\ \boldsymbol{\mu}^\top \beta &= 0 \\ \beta &\geq 0 \\ \gamma (\|\beta\|_p^p - 1) &= 0 \end{aligned} \quad (14)$$

where $\odot$ denotes the Hadamard (element-wise) multiplication. From the first KKT condition one obtains

$$|\beta|^{p-1} \odot \text{sign}(\beta) = \frac{\mathbf{u} + \boldsymbol{\mu}}{\gamma p} \quad (15)$$

Note that if $\gamma = 0$, then the p -norm constraint would not be satisfied and the optimisation problem would be unbounded.

Consequently, $\gamma > 0$, and thus, the equation above is well defined. At the optimum β is non-negative ($\beta \geq 0$) and hence

$$\beta = \left(\frac{\mathbf{u} + \boldsymbol{\mu}}{\gamma p} \right)^{1/(p-1)} \quad (16)$$

By plugging this into the second KKT condition one obtains

$$\boldsymbol{\mu}^\top \left[\frac{\mathbf{u} + \boldsymbol{\mu}}{\gamma p} \right]^{1/(p-1)} = 0 \quad (17)$$

Next, it is shown for $\boldsymbol{\mu}^\top \left[\frac{\mathbf{u} + \boldsymbol{\mu}}{\gamma p} \right]^{1/(p-1)} = 0$ to be satisfied, must be zero. To show $\boldsymbol{\mu} = 0$, contradiction is used. Let us suppose that all the elements of $\boldsymbol{\mu}$ are not zero and $\boldsymbol{\mu}$ includes one strictly positive element. Under this assumption, since γ , p and $\mathbf{u}$ are non-negative, the corresponding element in $\left[\frac{\mathbf{u} + \boldsymbol{\mu}}{\gamma p} \right]^{1/(p-1)}$ shall be strictly positive too. In this case, the inner product of $\boldsymbol{\mu}$ and β would be strictly positive which contradicts the initial requirement of $\boldsymbol{\mu}^\top \beta = 0$. Consequently, $\boldsymbol{\mu}$ will not include any strictly positive elements, and since it is non-negative, we conclude that $\boldsymbol{\mu} = 0$.

Inspecting the minimisation sub-problem w.r.t. β , it is evident that at the optimum, β_j 's should be as large as possible at the optimum which results in the maximisation of the p -norm. The maximum of the p -norm happens when $\gamma > 0$, requiring the p -norm constraint to hold as equality, and thus from the last KKT condition $\|\beta\|_p = 1$. Since $\gamma > 0$, $\boldsymbol{\mu} = 0$ and $\|\beta\|_p = 1$, β is derived as

$$\beta = \frac{\mathbf{u}^{1/(p-1)}}{\|\mathbf{u}^{1/(p-1)}\|_p} \quad (18)$$

The relation above is well defined for $1 < p < \infty$. The degenerate case of $p = 1$ and also the case of $p \rightarrow +\infty$ are analysed next.

$p \rightarrow 1^+$: Let us consider the numerator of the quotient in Eq. 18 in when $p \rightarrow 1^+$. Denoting the maximum element of $\mathbf{u}$ as u_{max} we have

$$\begin{aligned} \lim_{p \rightarrow 1^+} \mathbf{u}^{1/(p-1)} &= \lim_{p \rightarrow 1^+} [u_1, \dots, u_{max}, \dots, u_J]^{1/(p-1)} \\ &= \lim_{p \rightarrow 1^+} u_{max}^{1/(p-1)} \left[\left(\frac{u_1}{u_{max}} \right)^{1/(p-1)}, \dots, 1, \dots, \left(\frac{u_J}{u_{max}} \right)^{1/(p-1)} \right] \end{aligned} \quad (19)$$

Note that all elements of the vector above are smaller than one except for the element at the position of the maximum element which is equal to 1. Moreover, $\lim_{p \rightarrow 1^+} 1/(p-1) \rightarrow +\infty$ and consequently, in the limit when $p \rightarrow 1^+$, all the elements of the vector shrink to zero except for the maximum element. As a result, when $p \rightarrow 1^+$, β shall only have a single non-zero element of 1 at the position of the maximum element of $\mathbf{u}$.

$p \rightarrow +\infty$: In this case, we have $\lim_{p \rightarrow +\infty} 1/(p-1) = 0$. As a result, all the elements of $\mathbf{u}$ are raised to zero. Assuming that the elements of $\mathbf{u}$ are non-zero (if the kernel matrix is positive-definite this assumption holds true), β is then derived as a uniform vector with a unit p -norm.

Up to this point, a procedure to determine β is presented. Once β is determined, the objective function should be max-

Algorithm 1 ℓ_p -Norm Multiple Kernel One-Class Fisher Null-Space

```

1: Input: kernel matrices  $\mathbf{K}_j, j = 1, \dots, J$ 
2: Initialisation:  $\alpha = (\delta \mathbf{I} + \sum_{j=1}^J \frac{1}{j^{1/p}} \mathbf{K}_j)^{-1} \mathbf{1}$ 
3: repeat
4:    $\mathbf{u} = [\alpha^\top \mathbf{K}_1 \alpha, \dots, \alpha^\top \mathbf{K}_J \alpha]$ 
5:    $\beta = \frac{\mathbf{u}^{1/(p-1)}}{\|\mathbf{u}^{1/(p-1)}\|_p}$ 
6:    $\alpha = (\delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j)^{-1} \mathbf{1}$ 
7: until convergence
8: Output:  $\alpha$  and  $\beta$ 

```

imised in α . Setting the partial derivative of the cost function with respect to α to zero yields

$$\alpha = \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j \right)^{-1} \mathbf{1} \quad (20)$$

In the relation above, α is given in terms of β . If β was independent of α , then the optimal α could have been directly found. However, due to Eq. 18, β is given in terms of $\mathbf{u}$ which is a function of α . That is, in Eq. 20, α is defined as a function of itself which necessitates a different approach for optimisation. In this work, a fixed-point iteration procedure [42] is followed to solve for α . More specifically, let us define $f(\alpha) = \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j \right)^{-1} \mathbf{1}$, and hence at the optimum it must hold that $\alpha = f(\alpha)$. In order to determine α , we look for the fixed-point of function $f(\alpha)$. That is, at each iteration, using the current estimate of α , $f(\alpha)$ is determined which is then used as the new estimate for α . The procedure above is repeated until convergence. The approach described above is summarised in Algorithm 1 where α is initialised using a uniform β with a unit p -norm. A convergence analysis of the proposed approach is provided in Section III of the supplementary material.

A. Remark

The problem in Eq. 10 is concave in α . This may be readily observed by plugging β given in Eq. 18 into Eq. 12 to obtain (see Section IV in the supplementary material for a derivation)

$$\max_{\alpha} -\delta \alpha^\top \alpha + 2 \alpha^\top \mathbf{1} - \|\mathbf{u}\|_{p/(p-1)} \quad (21)$$

As the elements of $\mathbf{u}$ are quadratic functions of α , and since $\mathbf{K}_j$'s, $\forall j$ are positive semidefinite matrices, the elements of $\mathbf{u}$ are convex functions of α . Moreover, as any norm ($p \geq 1$) is a convex function, the problem above is concave. Consequently, one may consider a gradient ascent-type algorithm to solve for α . However, in practice, as will be demonstrated in Section VII-F, the fixed-point iteration approach presented in Algorithm 1 is significantly faster than a gradient-based method.

VI. EXTENSION TO JOINT ONE-CLASS MULTIPLE KERNEL LEARNING

In the discussions thus far, it was assumed that only a *single* one-class problem exists. In some applications, however, there might be several *related* one-class problems to be learned. In

Algorithm 2 Joint ℓ_p -Norm Multiple Kernel One-Class Fisher Null-Space

```

1: Input: kernel matrices  $\mathbf{K}_j^c, j = 1, \dots, J, c = 1, \dots, C$ 
2: Initialisation:
3: for  $c = 1, \dots, C$  do
4:    $\alpha_c = \left( \delta \mathbf{I} + \sum_{j=1}^J \frac{1}{j^{1/p}} \mathbf{K}_j^c \right)^{-1} \mathbf{1}_c$ 
5: end for
6: repeat
7:    $\mathbf{v} = [\sum_{c=1}^C \alpha_c^\top \mathbf{K}_1^c \alpha_c, \dots, \sum_{c=1}^C \alpha_c^\top \mathbf{K}_J^c \alpha_c]$ 
8:    $\beta = \frac{\mathbf{v}^{1/(p-1)}}{\|\mathbf{v}^{1/(p-1)}\|_p}$ 
9:   for  $c = 1, \dots, C$  do
10:     $\alpha_c = \left( \delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j^c \right)^{-1} \mathbf{1}_c$ 
11:   end for
12: until convergence
13: Output:  $\alpha_c$ 's and  $\beta$ 

```

such situations, different one-class tasks may share similarities that motivate inferring common kernel weights shared across all the OCC problems. This may be useful when the training samples for some one-class problems are scarce and not representative of the entire observation space. In this case, an MKL approach when applied separately to each individual one-class task may suffer from poor generalisation on unseen test data but taking all training observations from all one-class tasks may better represent the feature space. A further appealing characteristic of a joint one-class MKL approach is that of computational efficiency. In this context, under the assumption that different one-class MKL problems are related and share common kernel weights, once a joint learning is performed for a number of OCC MKL problems, any new and similar one-class MKL problem may then benefit from the weights previously inferred, avoiding the requirement for a further MKL learning procedure. In this section, an extension of Algorithm 1 to handle the case of joint multiple kernel learning for several related OCC problems is presented.

Let us assume that there exist C related one-class problems which are supposed to benefit from common kernel weights. We assume that each individual one-class problem (indexed by c) is specified by a distinct set of kernel matrices (i.e. $\mathbf{K}_j^c$, for $j = 1, \dots, J$) and a separate discriminant in the Hilbert space (i.e. α_c) but coupled to other OCC problems through a common kernel weight vector, i.e. β . With these assumptions, the optimisation problem for the joint ℓ_p -norm multiple kernel Fisher null-space may be written as

$$\min_{\beta} \max_{\alpha_c} \sum_{c=1}^C -\alpha_c^\top \left(\sum_{j=1}^J \beta_j \mathbf{K}_j^c \right) \alpha_c - \delta \alpha_c^\top \alpha_c + 2 \alpha_c^\top \mathbf{1}_c \quad (22)$$

s.t. $\beta \geq 0, \|\beta\|_p^p \leq 1$

where $\mathbf{1}_c$ is a vector of ones of a dimensionality suitable for the c^{th} one-class problem. If the kernel weights were learned independently for each individual OCC problem, the optimisation problem above would have been decoupled into C independent OCC MKL problems to each of which Algorithm 1 could be directly applied. However, the introduction of a common kernel weight vector β shared across several OCC tasks necessitates a modified optimisation procedure, discussed next. Let us define v_j to be the j^{th} element of the vector $\mathbf{v}$ given as $v_j = \sum_{c=1}^C \alpha_c^\top \mathbf{K}_j^c \alpha_c$. Applying the

minimax theorem, the saddle point problem in Eq. 22 may be written as

$$\begin{aligned} \max_{\alpha_c} \sum_{c=1}^C -\delta \alpha_c^\top \alpha_c + 2\alpha_c^\top \mathbf{1}_c + \min_{\beta} -\beta^\top \mathbf{v} \\ \text{s.t. } \beta \geq 0, \|\beta\|_p^p \leq 1 \end{aligned} \quad (23)$$

One notices the similarity of the minimisation subproblem above to that of Eq. 12. Hence, following a similar procedure, β is derived as

$$\beta = \frac{\mathbf{v}^{1/(p-1)}}{\|\mathbf{v}^{1/(p-1)}\|_p} \quad (24)$$

In order to maximise the objective function in Eq. 22 in α_c 's, one may take the partial derivatives with respect to each α_c and set them to zero to obtain

$$\alpha_c = \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j^c \right)^{-1} \mathbf{1}_c, \text{ for } c = 1, \dots, C \quad (25)$$

As in the single OCC scenario, a fixed-point algorithm can be deployed to find the optimal solution. The proposed joint multiple kernel one-class Fisher null-space approach is summarised as Algorithm 2 where α is initialised using a uniform unit p -norm kernel weight vector.

VII. EXPERIMENTAL EVALUATION

In this section, the results of an experimental evaluation of the proposed approach and a comparison against other methods from the literature are provided. The kernel-based methods included in the comparison are:

- *FN-Ave.*: As a first baseline method, this method corresponds to the element-wise arithmetic mean of the kernels followed by the kernel Fisher null-space approach for OCC [17], [18];
- *FN-Prod.*: As a second baseline method, this method corresponds to the element-wise geometric mean of the kernels followed by the kernel Fisher null-space approach for OCC [17], [18];
- *GP-Ave.*: corresponds to the element-wise arithmetic mean of the kernels followed by the kernel GP method for OCC [15];
- *GP-Prod.*: corresponds to the element-wise geometric mean of the kernels followed by the kernel GP method for OCC [15];
- *KPCA-Ave.*: corresponds to the element-wise arithmetic mean of the kernels followed by the KPCA method for OCC [16];
- *KPCA-Prod.*: corresponds to the element-wise geometric mean of the kernels followed by the KPCA method for OCC [16];
- *MK-SVDD*: is the multiple kernel learning method based on the SVDD approach for OCC [36];
- *Slim-MK-SVDD*: is the multiple kernel learning method based on the SVDD approach for OCC where the objective function is modified so that tight solutions are favoured over loose ones [36];

- *MK-OCSVM*: is the multiple kernel learning method based on the One-Class SVM algorithm for OCC [35];
- *Slim-MK-OCSVM*: is the multiple kernel learning method based on the One-Class SVM algorithm for OCC where the objective function is modified so that tight solutions are favoured over loose ones [36];
- ℓ_p *MK-FN*: is the proposed ℓ_p -norm multiple kernel Fisher null-space approach;
- *Joint ℓ_p MK-FN*: is the proposed ℓ_p -norm multiple kernel Fisher null-space approach where a shared kernel weight is inferred for several related one-class problems.

In addition to the aforementioned kernel-based methods, wherever applicable, we incorporate state-of-the-art end-to-end learning deep one-class methods in the experiments for a comparison.

The remainder of the experimental evaluation section is organised as detailed next.

- Section VII-A provides the details of implementation.
- In Section VII-B, we analyse the convergence characteristics of the proposed method.
- In Section VII-C, the proposed approach is evaluated for abnormality detection and compared against the baseline as well as other one-class multiple kernel fusion and deep end-to-end OCC methods.
- In Section VII-D, the proposed approach is evaluated for one-class novelty detection and compared to the baseline as well as other one-class multiple kernel fusion and deep end-to-end OCC methods.
- In Section VII-E, we conduct experiments in biometric anti-spoofing (presentation attack detection) and compare the performance of the proposed approach to the baseline as well as other one-class multiple kernel fusion and deep end-to-end OCC methods.
- Finally, in Section VII-F, we compare the running time of different one-class MKL methods and discuss the scalability properties of the proposed approach.

A. Implementation details

Motivated by earlier studies in the multi-class setting, the parameter p in the proposed approach is chosen from $\{1, 32/31, 16/15, 8/7, 4/3, 2, 4, 8, 10^6\}$. The regularisation parameter δ is selected from $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 10^2\} \times n$. The kernel function used in the experiments is that of a Gaussian (RBF) kernel the width of which is chosen from $\{\frac{1}{4}D, \frac{1}{2}D, D\}$ where D is the average pairwise Euclidean distance between training samples. The parameters in the proposed approach are determined on the validation set. The parameters of the multiple kernel learning methods based on SVM are set as suggested in [36] on the validation set.

B. Convergence Behaviour

In this section, we examine the convergence behaviour of the proposed ℓ_p -norm one-class MKL approach. For this purpose, we randomly choose a single class from the Oxford Flowers17 [43] dataset (a dataset of flower types) and identify

Fig. 2. Convergence curves for a sample multiple kernel learning one-class problem with regards to different regularisations in terms of parameter p (see Section VII-B for further details).

it as the target class. Seven distances matrices based on different colour spaces, texture and shape features of flower images are available online¹. Based on the provided distance matrices, 7 kernel matrices are built. Using the proposed approach in Algorithm 1 ($\delta = 10^{-3}n$), we then estimate α and the optimal kernel weights. As a measure of convergence we record the l_2 -norm of the change in α . A zero change is indicative of convergence. The experiment is repeated 100 times for each value of p where at each iteration we randomly initialise β to a non-negative vector with a unit p -norm. The results are depicted in Fig. 2 where the vertical axis denoted as "Error" represents the l_2 -norm of the change in α . The solid curve depicts the average while the shaded region represents the standard deviation of the error. From the figure, it can be seen that the convergence is quite fast. More specifically, for smaller values of p (i.e. for $p = 32/31, 16/15, 8/7$ and $4/3$) convergence is typically attained in as few as 10 iterations. For larger p 's, the convergence is even faster where for $p = 2, 4, 8$ and 10^6 , in most cases, a maximum of 3 iterations suffices for convergence.

C. Abnormal Image Detection

In abnormal image detection the task is to label images as normal or abnormal. As the form of the abnormality in the images is not known a priori, training is performed on images corresponding to the normal/positive class. One of the dedicated datasets for visual abnormality is the 1001 Abnormal Objects database [44] that incorporates 1001 abnormal images from six classes of object that originally appeared in the PASCAL database [45]. The groundtruth labels regarding abnormality of an object in this dataset are obtained using human responses collected from the Amazon Mechanical Turk. The 6 object classes included in the 1001 Abnormal Objects dataset are Boat, Airplane, Chair, Car, Sofa and Motor-bike where the number of abnormal images associated with each object class

TABLE I
COMPARISON OF DIFFERENT METHODS FOR DETECTING ABNORMAL IMAGES ON THE ABNORMALITY-1001 DATABASE.

Method	AUC (mean $\pm$ std %)
FN-Ave.	94.7 $\pm$ 0.2
FN-Prod.	90.1 $\pm$ 0.3
GP-Ave.	93.5 $\pm$ 0.2
GP-Prod.	84.4 $\pm$ 0.3
KPCA-Ave.	91.6 $\pm$ 0.2
KPCA-Prod.	87.3 $\pm$ 0.3
MK-SVDD	92.3 $\pm$ 0.2
Slim-MK-SVDD	94.1 $\pm$ 0.1
MK-OCSVM	91.4 $\pm$ 0.3
Slim-MK-OCSVM	93.9 $\pm$ 0.1
Graphical Model [44]	87.0 $\pm$ <i>n.a.</i>
Adjusted Graphical Model [44]	91.1 $\pm$ <i>n.a.</i>
Autoencoder [54]	67.4 $\pm$ 1.2
OCNN [55]	88.5 $\pm$ 1.4
DOC [53]	95.6 $\pm$ 3.1
OC-CNN [56]	84.3 $\pm$ <i>n.a.</i>
ℓ_p MK-FN	96.2 $\pm$ 0.3

in the dataset is at least 100. We construct seven kernels for this dataset using deep representations obtained from the pre-trained deep CNN's of Googlenet [46], Resnet50 [47], Vgg16 [48], Alexnet [49], Nasnetlarge [50], Mobilenetv2 [51] and Densenet201 [52]. We follow the protocol associated with this dataset introduced in [44] to allow for a fair comparison with other methods in the literature. For the comparison with the proposed one-class ℓ_p MK-FN approach, we include ten other multiple kernel fusion methods introduced earlier as well as state-of-the-art approaches from the literature including end-to-end deep one-class learning methods. The performances of different methods on this dataset are reported in Table I. From the table, one may observe that the proposed ℓ_p MK-FN approach performs better than other multiple kernel methods. The proposed approach achieves an average AUC of 96.2% whereas the second best performing kernel-based method obtains an average AUC of 94.7%. It is interesting to note that the second best performing multiple kernel method does not correspond to one of the SVM-based MKL methods but to a simple average fusion of multiple kernels in the Fisher null-space framework. The best performance in this case corresponds to the one-class deep end-to-end method in [53] where 95.6% for an average AUC is obtained which is worst compared to the performance of the proposed approach which underlines the utility of an optimal combination of multiple kernels (as in the proposed approach) associated with different deep representations for visual abnormality detection.

D. One-Class Novelty Detection

The goal in one-class novelty detection is to gauge the novelty of a new observation using the previously enrolled data items. As the type of novelty in test observations is unknown before the testing stage, training may be performed using one-class classification strategies. Two of the widely used datasets to evaluate the performance of different novelty detection methods are those of the Caltech 256 dataset [57] and the MNIST dataset [58]. The Caltech 256 dataset incorporates object images from 256 different classes for a total of 30607 object images. For the purpose of conducting a fair comparison

¹<http://www.robots.ox.ac.uk/~vgg/data/flowers/17/index.html>

to other methods, we follow two different protocols. The first protocol is the one used in [59] which considers each single class as the target class and assumes samples from all the other 255 object classes as novelties. The experiment is repeated in turn for the first 40 classes in the dataset and the performance is reported in terms of the AUC. The second protocol followed for the Caltech256 dataset is the one employed in [60] which considers each one of five randomly chosen categories as the positive class and the observations from the "clutter" category as outliers. The second dataset used for novelty detection in this study is that of MNIST [58] which is comprised of 60,000 handwritten digits from "0" to "9". For the purpose of this experiment, each digit is regarded as the target class while all other digits serve as novelties. We repeat the experiment for all digits in turn. In this experiment, we report the performances in terms of average AUC over all classes. Multiple kernels are constructed for the images in both datasets using the seven pre-trained deep CNN models similar to the previous experiment.

Similar to the Abnormal-1001 dataset, for the comparison with the proposed approach, ten other multiple kernel fusion methods as well as state-of-the-art methods including one-class end-to-end deep methods are considered. The results of this experiment are reported in Table II and Table III for the first and the second protocol on the Caltech256 dataset, respectively. Table IV reports the evaluation results on the MNIST dataset. A couple of observations from these experiments are in order. First, in both evaluation settings on the Caltech256 dataset, all the MK approaches perform better than other methods including one-class deep methods of [55], [53], [54], [53] as well as non-deep methods such as R-graph [61]. Among the multiple kernel methods, similar to the abnormal image dataset, while the method with the best performance is the proposed ℓ_p -norm one-class MKL approach, the second best performing method does not correspond to any of the SVM-based MKL approaches but to other fixed-rule multiple kernel methods of the Fisher null-space, GP and the KPCA approaches. The best performing method in both protocols of the Caltech256 dataset is that of the proposed ℓ_p MK-FN approach with an average AUC of 99.5% for the 40-class scenario and an average AUC of 99.8 for the 5-class setting. The best performing deep one-class learning method from the literature for the 40-class setting corresponds to the method in [53] with an average AUC of 98.1% while the best performing method from the literature for the 5-category case is that of [60] with an average AUC of 92.3%. On the MNIST dataset, the proposed ℓ_p MK-FN approach obtains the best performance with an average AUC of 98.7% while the best performing methods from the literature correspond to those in [62] and [63] with an average AUC of 97.5%. Similar to the Caltech256 and the Abnormality-1001 datasets, the Fisher null space approach even when using a fixed kernel fusion rule provides very competitive performance to the-state-of-the-art methods which illustrates the efficacy of the Fisher null classification principle for OCC.

E. Face Presentation Attack (Spoofing) Detection

In face biometrics, an important application of one-class classification is that of presentation attack detection (PAD)

TABLE II
COMPARING DIFFERENT METHODS FOR DETECTING NOVELTIES ON THE CALTECH 256 DATABASE USING THE FIRST PROTOCOL.

Method	AUC (mean $\pm$ std %)
FN-Ave.	99.3 $\pm$ 0.6
FN-Prod.	99.3 $\pm$ 0.6
GP-Ave.	99.3 $\pm$ 0.6
GP-Prod.	99.3 $\pm$ 0.6
KPCA-Ave.	99.2 $\pm$ 0.8
KPCA-Prod.	99.3 $\pm$ 0.8
MK-SVDD	99.1 $\pm$ 0.7
Slim-MK-SVDD	99.2 $\pm$ 0.6
MK-OCSVM	99.1 $\pm$ 0.8
Slim-MK-OCSVM	99.1 $\pm$ 0.8
OCNN VGG16 [55]	88.5 $\pm$ 14.4
OCNN AlexNet [55]	82.6 $\pm$ 15.3
OCSVM VGG16 [53]	90.2 $\pm$ 5.0
Autoencoder [54]	62.3 $\pm$ 7.2
DOC AlexNet [53]	95.9 $\pm$ 2.1
DOC VGG16 [53]	98.1 $\pm$ 2.2
ℓ_p MK-FN	99.5 $\pm$ 0.5

TABLE III
COMPARING DIFFERENT METHODS FOR DETECTING NOVELTIES ON THE CALTECH 256 DATABASE USING THE SECOND PROTOCOL.

Method	AUC (mean $\pm$ std %)
FN-Ave.	99.6 $\pm$ 0.4
FN-Prod.	99.5 $\pm$ 0.3
GP-Ave.	99.7 $\pm$ 0.3
GP-Prod.	99.6 $\pm$ 0.4
KPCA-Ave.	99.4 $\pm$ 0.7
KPCA-Prod.	99.4 $\pm$ 0.6
MK-SVDD	99.4 $\pm$ 0.6
Slim-MK-SVDD	99.4 $\pm$ 0.4
MK-OCSVM	99.3 $\pm$ 0.5
Slim-MK-OCSVM	99.3 $\pm$ 0.7
DPCP [64]	67.7 $\pm$ <i>n.a.</i>
R-graph [61]	91.3 $\pm$ <i>n.a.</i>
The work in [60]	92.3 $\pm$ <i>n.a.</i>
ℓ_p MK-FN	99.8 $\pm$ 0.2

where the goal is to determine whether a presentation made to the system sensor corresponds to a genuine (bona fide) biometric trait or is a reproduction of a subject's biometric data (presentation attack) to gain illegitimate access to the system. Typical examples of presentation attacks include printed photo attack and video display attack. As new and previously unseen attacks in the training set may potentially be developed by attackers, the problem is quite a challenging one, emphasising the need for a system with high generalisation capability to detect unseen attacks. One of the promising techniques to deal with unseen attacks is known to be that of one-class

TABLE IV
COMPARING DIFFERENT METHODS FOR DETECTING NOVELTIES ON THE MNIST DATABASE.

Method	AUC %
FN-Ave.	97.9
FN-Prod.	93.2
GP-Ave.	96.3
GP-Prod.	87.7
KPCA-Ave.	94.9
KPCA-Prod.	90.2
MK-SVDD	95.7
Slim-MK-SVDD	97.4
MK-OCSVM	94.6
Slim-MK-OCSVM	97.1
The work in [63]	97.5
OCGAN [62]	97.5
the work in [65]	97.1
ℓ_p MK-FN	98.7

classification [4]. In this case, bona fide samples are considered as target objects and presentation attacks (PA's) as anomalies. The face presentation attack detection (PAD) problem may be approached in a subject-specific fashion by training a distinct one-class learner for each individual subject enrolled in the dataset [66]. In this work, the 12 kernels developed in [67] are utilised. The kernels correspond to the deep features extracted by applying the pre-trained GoogleNet [46], ResNet50 [47] and VGG16 [48] networks onto four different facial regions including the whole face as well as different local facial regions that correspond to the main facial features including the joint nose and the eyes region, the nose region and the lower part of the face. We conduct the experiment in an unseen scenario meaning that only positive (i.e. bona fide) samples are utilised for training. The standard metrics for performance reporting in biometric PAD are ISO metrics BSISO-IEC30107-3-2017 [68] defined as 1) attack presentation classification error rate (APCER for short) that represents the percentage of attack presentation attempts via the same PA instrument species which are misclassified as bona fide observations; and 2) bona fide presentation classification error rate (BPCER for short) that corresponds to the proportion of bona fide presentation attempts that are misclassified as PA's. The performance of a PAD system in detecting PA samples is reported using the highest APCER among all presentation attack instrument species (PAIS):

$$APCER = \max_{PAIS} APCER_{PAIS} \quad (26)$$

The overall performance of the PAD system may be summarised in terms of the Average Classification Error Rate (ACER):

$$ACER = \frac{\max_{PAIS} APCER_{PAIS} + BPCER}{2} \quad (27)$$

The face presentation attack dataset used in this study is the OULU-NPU dataset [69] which is comprised of 4950 bona fide videos and attack samples from 55 subjects. The samples are partitioned into three subsets of non-overlapping individuals for training (360 bona fide videos and 1440 PA items), development (270 bona fide samples and 1080 PA videos) and testing (360 bona fide samples and 1440 PA videos). Among the four evaluation protocols of the dataset, the most challenging one is the forth protocol which is used in this study to simultaneously evaluate the performance across previously unseen illumination conditions, background scenes, PAIs and input sensors. The evaluation protocol requires to report the average and the standard deviation of the performance over six different mobile devices used to collect the dataset.

The Oulu-NPU dataset lack subject-specific validation data. As such, for the " ℓ_p MK-FN" approach, for each subject in the test set, we use the data corresponding to all subjects but the subject of interest as the validation data. For the "Joint ℓ_p MK-FN" approach, we use the data associated with the subjects other than the test subjects as the validation set. Apparently, this is not an ideal choice to tune system parameters which is forced by the limitations in terms of data and can potentially lead to a sub-optimal performance. As the one-class problems in this dataset share similarities (i.e. they all correspond to face

Fig. 3. Shared kernel weights inferred for the Oulu-NPU dataset for different p -norm regularisations.

TABLE V
COMPARING DIFFERENT METHODS FOR THE UNSEEN ZERO-SHOT FACE SPOOFING ON THE OULU-NPU DATABASE PROTOCOL IV.

Method	ACER (mean $\pm$ std %)
FN-Ave.	5.0 $\pm$ 3.9
FN-Prod.	4.5 $\pm$ 5.3
GP-Ave.	6.2 $\pm$ 4.4
GP-Prod.	5.8 $\pm$ 6.4
KPCA-Ave.	5.4 $\pm$ 3.6
KPCA-Prod.	4.5 $\pm$ 5.3
MK-SVDD	7.1 $\pm$ 6.2
Slim-MK-SVDD	6.2 $\pm$ 4.4
MK-OCSVM	7.9 $\pm$ 6.4
Slim-MK-OCSVM	6.2 $\pm$ 4.4
OCA-FAS [70]	4.1 $\pm$ 2.7
the work in [71]	3.7 $\pm$ 2.1
the work in [72]	9.8 $\pm$ 4.2
SAPLC [73]	9.3 $\pm$ 4.4
ℓ_p MK-FN	3.3 $\pm$ 3.4
Joint ℓ_p MK-FN	3.3 $\pm$ 3.0

images) one may infer common kernel weights shared across all one-class problems. The shared kernel weights for the test subjects in this dataset are depicted in Fig. 3 and the ACER's are reported in Table V where ten multiple kernel methods as well as the best performing end-to-end deep methods from the literature are included for a comparison. From Table V, the following observations can be made. First, the fixed-rule one-class Fisher-based and the KPCA-based methods perform better than the SVM-based MKL approaches. Among the SVM-based MK methods, the "slim" variants of the SVDD and the OCSVM methods appear to perform better than their naïve versions, as expected. Among the multiple kernel methods, the proposed " ℓ_p MK-FN" approach achieves the best performance. Moreover, it outperforms the best reported unseen PAD performance on this dataset corresponding to the method in [71] based on deep end-to-end learning for anomaly detection. The "Joint ℓ_p MK-FN" method appears to perform relatively well. In particular, while it archives a similar average error rate as that of the ℓ_p MK-FN method, nevertheless, in practical settings, it provides an additional advantage to the ℓ_p MK-FN in terms of computational complexity as any new user may directly benefit from the previously learned kernel weights avoiding the requirement for a new subject-specific multiple kernel learning stage.

F. Running Time of Different One-Class MKL Methods

In this section, a comparison of different one-class MKL methods in terms of their running times is provided. Two

TABLE VI

AVERAGE TRAINING TIMES OF DIFFERENT ONE-CLASS MULTIPLE KERNEL LEARNING METHODS (IN MILLISECONDS) FOR DIFFERENT NUMBER OF KERNELS (J). (THE NUMBER OF TRAINING SAMPLES IS 100, THEIR DIMENSIONALITY IS 100.)

Number of kernels	5	15	50	100	500
ℓ_p MK-FN	1.8	4.2	13.9	4.2	33.4
ℓ_p MK-FN-Grad	35.4	410.5	420.5	4893.2	115881.3
ℓ_p MK-FN-CVX	334.5	321.4	346.4	516.0	1177.9
Slim-MK-SVDD	34.1	65.3	296.1	589.1	1696.0
MK-SVDD	78.5	173.7	325.6	1228.9	4094.6
Slim-MK-OCSVM	41.9	93.7	221.4	982.1	685.8
MK-OCSVM	49.2	85.4	324.7	1293.4	2608.9

TABLE VII

AVERAGE TRAINING TIMES OF DIFFERENT ONE-CLASS MULTIPLE KERNEL LEARNING METHODS (IN MILLISECONDS) FOR DIFFERENT NUMBER OF TRAINING SAMPLES (n). (THE DIMENSIONALITY OF TRAINING SAMPLES IS 100 AND THE NUMBER OF KERNELS IS SET TO 10.)

No. of training samples	100	1000	5000	10000
ℓ_p MK-FN	2.1	719.0	6820.8	21315.7
ℓ_p MK-FN-Grad	58.3	20713.0	135635.4	323040.8
ℓ_p MK-FN-CVX	529.2	1066.8	10387.3	37022.8
Slim-MK-SVDD	605.7	30392.2	85023.3	497251.3
MK-SVDD	63.5	2172.2	26094.2	108472.1
Slim-MK-OCSVM	404.0	39973.4	401301.7	1753212.8
MK-OCSVM	46.7	4254.1	55474.9	163852.4

sets of experiments are conducted to study the impact of the number of kernels and the number of training samples on training times. Note that as discussed previously, the optimisation problem in Eq. 11 is concave, and hence, it may be solved via convex optimisation tools. As an additional method, we solve the minimisation subproblem in Eq. 11 using CVX, a package for specifying and solving convex programmes [74] and refer to it as ℓ_p MK-FN-CVX in the comparisons. Moreover, we examine a gradient ascent-based approach for optimisation which is referred to as ℓ_p MK-FN-Grad (please consult Section V of the supplementary material for a derivation of the gradient of the cost function). Note that as we set the stopping condition for the ℓ_p MK-FN, ℓ_p MK-FN-CVX and the ℓ_p MK-FN-Grad methods similarly, their results exactly match, and hence, their running times are comparable. We generate the data for this experiment randomly (as discussed in Section VI.A of the supplementary material) and set p to 2. The proposed method is implemented as un-optimised Matlab codes. The implementation of the SVM-based MKL methods are due to [36]. Each experiment is repeated 100 times on a 64-bit 4.00GHz Intel Core-i7 machine with 32GB memory and we report the average training times over 10 trials. The results of this experiment are reported in Table VI and Table VII. In order to better compare the running time of different one-class MKL algorithms in the aforementioned tables we have excluded the time required to compute the kernel matrix which is common to all algorithms and has a time complexity of $\mathcal{O}(n^2)$ on a serial machine. Nevertheless, kernel matrix computation may benefit from a parallel processing hardware such as a GPU to achieve large speed-up gains.

From the tables the following observations are in order. The proposed ℓ_p MK-FN approach when compared to the

ℓ_p MK-FN-CVX and the MK-FN-Grad is faster by multiple orders of magnitude. When compared to the SVM-based MKL methods, the proposed ℓ_p MK-FN method also comes out as the most efficient method in the majority of the experiments. In summary, the proposed ℓ_p MK-FN method outperforms other MKL methods in different scenarios considered above. The work in [17] presents an efficient one-class Fisher null classifier via an incremental approach. However, due to the absence of a Tikhonov regularisation on the solution, it may not be directly plugged into the proposed approach since the convergence of the proposed MKL algorithm (discussed in Section III of the supplementary material) heavily relies on the Tikhonov regularisation. An investigation on the modifications required to be applied on the algorithm presented in [17] so that it could be employed within the proposed MKL framework may be considered as a future direction of investigation.

1) *Scalability*: The computational complexity of the proposed method is dominated by the computation of vector $\mathbf{u}$ requiring matrix-vector multiplications (step 4 of Algorithm 1), the time complexity of which is $\mathcal{O}(Jn^2)$ where J is the number of kernels and n denotes the number of training samples. Note that performing a naïve matrix inversion operation (step 6 of Algorithm 1) would require $\mathcal{O}(n^3)$ time. However, the computational complexity of matrix inversion may be reduced to $\mathcal{O}(n^2)$ using an incremental Cholesky decomposition [75], [76]. An appealing characteristic of the proposed method is that all the operations involved are parallelizable. In particular, the computation of vector $\mathbf{u}$ (step 4 of Algorithm 1) that involves matrix-vector multiplications can be readily ported onto parallel processing units such as GPU's to obtain massive speed-up gains. Similarly, there has been a large volume of work on parallel implementations of the matrix inversion operation where it is shown that significant improvements in the running time may be obtained [77], [78].

The parallelizability attribute discussed above also applies to the *Joint* ℓ_p -norm MKL approach. Moreover, the joint approach provides a particularly appealing additional property which may be deployed for reduced time complexity: under the assumption that the one-class MKL problems share similarities in terms of feature space representations, one may learn the *shared* kernel combination weights for a number of OCC problems offline and directly apply the inferred weights to a new MKL OCC problem, circumventing the computational overhead corresponding to the new OCC MKL problem. As an instance, in the face PAD problem, using a number of subjects previously enrolled into the system, the optimal kernel combination weights may be inferred using the *Joint* ℓ_p -norm MKL algorithm. A new subject to be registered to the system may then benefit from the common kernel combination weights previously learned, completely avoiding any new one-class multiple kernel learning procedure.

VIII. CONCLUSION

We addressed the MKL problem for one-class classification. For this purpose, based on the one-class Fisher null approach, an ℓ_p -norm one-class MKL method was presented. The associated optimisation problem was posed as a saddle

point Lagrangian optimisation problem which was then solved via a new fixed-point iteration algorithm. The proposed approach was also extended to jointly learn multiple related one-class problems by constraining them to share common kernel weights. We theoretically studied the extreme cases of the proposed approach and illustrated that the ℓ_1 -norm only selects a single kernel from among a set of candidate kernels while the ℓ_∞ yields uniform kernel weights.

The evaluation of the proposed method on a range of datasets from different application domains illustrated the merits of the proposed method against other alternatives. In particular, it was shown that a fixed-norm solution may be outperformed by learning the intrinsic sparsity of the problem at hand as considered in the proposed approach. The proposed ℓ_p MK-FN approach was also demonstrated to perform better than the one-class SVM-based MKL methods as well as state-of-the-art one-class end-to-end deep methods from the literature. As future directions of investigation one may consider a mixed (r, p) -norm regularisation scheme on kernel weights (similar in spirit to the work in [33]) to provide further flexibility to the one-class MKL model.

REFERENCES

- [1] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 15:1–15:58, Jul. 2009.
- [2] S. S. Khan and M. G. Madden, "One-class classification: taxonomy of study and review of techniques," *The Knowledge Engineering Review*, vol. 29, no. 3, p. 345–374, 2014.
- [3] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, pp. 1263–1284, Sep. 2009.
- [4] S. R. Arashloo, J. Kittler, and W. Christmas, "An anomaly detection approach to face spoofing detection: A new formulation and evaluation protocol," *IEEE Access*, vol. 5, pp. 13 868–13 882, 2017.
- [5] M. Yu, Y. Yu, A. Rhuma, S. M. R. Naqvi, L. Wang, and J. A. Chambers, "An online one class support vector machine-based person-specific fall detection system for monitoring an elderly individual in a room environment," *IEEE Journal of Biomedical and Health Informatics*, vol. 17, no. 6, pp. 1002–1014, Nov 2013.
- [6] A. Rabaoui, M. Davy, S. Rossignol, and N. Ellouze, "Using one-class svms and wavelets for audio surveillance," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 4, pp. 763–775, Dec 2008.
- [7] P. Nader, P. Honeine, and P. Beausery, "lp-norms in one-class classification for intrusion detection in scada systems," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2308–2317, Nov 2014.
- [8] S. Budalakoti, A. N. Srivastava, and M. E. Otey, "Anomaly detection and diagnosis algorithms for discrete symbol sequences with applications to airline safety," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 39, no. 1, pp. 101–113, Jan 2009.
- [9] S. Kamaruddin and V. Ravi, "Credit card fraud detection using big data analytics: Use of psaoann based one-class classification," in *Proceedings of the International Conference on Informatics and Analytics*, ser. ICIA-16. New York, NY, USA: ACM, 2016, pp. 33:1–33:8. [Online]. Available: <http://doi.acm.org/10.1145/2980258.2980319>
- [10] G. G. Sundarkumar, V. Ravi, and V. Siddeshwar, "One-class support vector machine based undersampling: Application to churn prediction and insurance fraud detection," in *2015 IEEE International Conference on Computational Intelligence and Computing Research (ICICR)*, Dec 2015, pp. 1–7.
- [11] D. Tax, "One-class classification: concept-learning in the absence of counter-examples," Ph.D. dissertation, Delft University of Technology, 2001, aSCI Dissertation Series 65.
- [12] D. M. Tax and R. P. Duin, "Support vector data description," *Machine Learning*, vol. 54, no. 1, pp. 45–66, Jan 2004. [Online]. Available: <https://doi.org/10.1023/B:MACH.0000008084.60811.49>
- [13] H. Jiang, H. Wang, W. Hu, D. Kakde, and A. Chaudhuri, "Fast incremental svdd learning algorithm with the gaussian kernel," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, pp. 3991–3998, Jul. 2019. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/4291>
- [14] B. Schölkopf, J. C. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Comput.*, vol. 13, no. 7, pp. 1443–1471, Jul. 2001. [Online]. Available: <https://doi.org/10.1162/089976601750264965>
- [15] M. Kemmler, E. Rodner, E.-S. Wacker, and J. Denzler, "One-class classification with gaussian processes," *Pattern Recognition*, vol. 46, no. 12, pp. 3507 – 3518, 2013.
- [16] H. Hoffmann, "Kernel pca for novelty detection," *Pattern Recognition*, vol. 40, no. 3, pp. 863 – 874, 2007.
- [17] J. Liu, Z. Lian, Y. Wang, and J. Xiao, "Incremental kernel null space discriminant analysis for novelty detection," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 4123–4131.
- [18] P. Bodesheim, A. Freytag, E. Rodner, M. Kemmler, and J. Denzler, "Kernel null space methods for novelty detection," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, June 2013, pp. 3374–3381.
- [19] S. R. Arashloo and J. Kittler, "One-class kernel spectral regression for outlier detection," *CoRR*, vol. abs/1807.01085, 2018. [Online]. Available: <http://arxiv.org/abs/1807.01085>
- [20] X. Huang, J. Xu, and G. Guo, "Incremental kernel null foley-sammon transform for person re-identification," in *2018 24th International Conference on Pattern Recognition (ICPR)*, Aug 2018, pp. 1683–1688.
- [21] S. R. Arashloo and J. Kittler, "Robust one-class kernel spectral regression," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 3, pp. 999–1013, 2021.
- [22] M. Gönen and E. Alpaydin, "Multiple kernel learning algorithms," *Journal of Machine Learning Research*, vol. 12, no. 64, pp. 2211–2268, 2011. [Online]. Available: <http://jmlr.org/papers/v12/gonen11a.html>
- [23] Q. Mao, I. W. Tsang, S. Gao, and L. Wang, "Generalized multiple kernel learning with data-dependent priors," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 6, pp. 1134–1148, 2015.
- [24] Y. Han, K. Yang, Y. Yang, and Y. Ma, "Localized multiple kernel learning with dynamical clustering and matrix regularization," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 2, pp. 486–499, 2018.
- [25] S. S. Bucak, R. Jin, and A. K. Jain, "Multiple kernel learning for visual object recognition: A review," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 7, pp. 1354–1369, 2014.
- [26] H. Xia, S. C. H. Hoi, R. Jin, and P. Zhao, "Online multiple kernel similarity learning for visual search," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 536–549, 2014.
- [27] Y. Lin, T. Liu, and C. Fuh, "Multiple kernel learning for dimensionality reduction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 6, pp. 1147–1160, 2011.
- [28] X. Liu, L. Wang, X. Zhu, M. Li, E. Zhu, T. Liu, L. Liu, Y. Dou, and J. Yin, "Absent multiple kernel learning algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 6, pp. 1303–1316, 2020.
- [29] N. Subrahmanya and Y. C. Shin, "Sparse multiple kernel learning for signal processing applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 5, pp. 788–798, 2010.
- [30] F. Yan, J. Kittler, K. Mikołajczyk, and A. Tahir, "Non-sparse multiple kernel fisher discriminant analysis," *Journal of Machine Learning Research*, vol. 13, no. 21, pp. 607–642, 2012. [Online]. Available: <http://jmlr.org/papers/v13/yan12a.html>
- [31] M. Kloft, U. Brefeld, S. Sonnenburg, and A. Zien, "lp-norm multiple kernel learning," *Journal of Machine Learning Research*, vol. 12, no. 26, pp. 953–997, 2011. [Online]. Available: <http://jmlr.org/papers/v12/kloft11a.html>
- [32] C. Cortes, M. Mohri, and A. Rostamizadeh, "L2 regularization for learning kernels," in *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI 2009)*, Montréal, Canada, 2009. [Online]. Available: <http://www.cs.nyu.edu/~mohri/postscript/l2reg-uai.pdf>
- [33] Y. Han, Y. Yang, X. Li, Q. Liu, and Y. Ma, "Matrix-regularized multiple kernel learning via (r, p) norms," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 10, pp. 4997–5007, 2018.
- [34] G. R. Lanckriet, N. Cristianini, P. Bartlett, L. E. Ghaoui, and M. I. Jordan, "Learning the kernel matrix with semidefinite programming," *Journal of Machine Learning Research*, vol. 5, no. Jan, pp. 27–72, 2004. [Online]. Available: <http://www.jmlr.org/papers/v5/lanckriet04a.html>

- [35] A. Rakotomamonjy, F. Bach, S. Canu, Y. Grandvalet *et al.*, “Simpleml,” *Journal of Machine Learning Research*, vol. 9, pp. 2491–2521, 2008.
- [36] G. Loosli and H. Aboubacar, “Using SVDD in SimpleMKL for 3D-Shapes Filtering,” in *CAP 2016*, Marseille, France, Jul. 2016. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01593595>
- [37] S. Sonnenburg, G. Rätsch, C. Schäfer, and B. Schölkopf, “Large scale multiple kernel learning,” *Journal of Machine Learning Research*, vol. 7, pp. 1531–1565, Jul. 2006.
- [38] M. Gönen and E. Alpaydın, “Localized algorithms for multiple kernel learning,” *Pattern Recognition*, vol. 46, no. 3, pp. 795 – 807, 2013.
- [39] K. Hlaváčková-Schindler, “Tikhonov regularization parameter in reproducing kernel hilbert spaces with respect to the sensitivity of the solution,” in *Artificial Neural Networks - ICANN 2008*, V. Kůrková, R. Neruda, and J. Koutník, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 215–224.
- [40] D. Du and P. Pardalos, *Minimax and applications*, ser. Nonconvex Optimization and Its Applications. Kluwer Academic Publishers, 1995.
- [41] J. von Neumann, “Über ein ökonomisches gleichungssystem und eine verallgemeinerung des brouwerschen fixpunktsatzes,” *Wein 8 Ergebn. Math. Kolloq. Wein*, pp. 73–83, 1937.
- [42] J. Hoffman and S. Frankel, *Numerical Methods for Engineers and Scientists*, 2nd ed. CRC Press, 2001.
- [43] M. Nilsback and A. Zisserman, “Automated flower classification over a large number of classes,” in *2008 Sixth Indian Conference on Computer Vision, Graphics Image Processing*, 2008, pp. 722–729.
- [44] B. Saleh, A. Farhadi, and A. Elgammal, “Object-centric anomaly detection by attribute-based reasoning,” in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 787–794.
- [45] M. Everingham, L. V. Gool, C. K. I. Williams, J. Winn, and A. Zisserman, “The pascal visual object classes (voc) challenge,” *International Journal of Computer Vision*, vol. 88, pp. 303–308, September 2009, printed version publication date: June 2010.
- [46] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 1–9.
- [47] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [48] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2014.
- [49] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012, pp. 1097–1105.
- [50] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, “Learning transferable architectures for scalable image recognition,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8697–8710.
- [51] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, “Mobilenetv2: Inverted residuals and linear bottlenecks,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520.
- [52] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely connected convolutional networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269.
- [53] P. Perera and V. M. Patel, “Learning deep features for one-class classification,” *IEEE Transactions on Image Processing*, vol. 28, no. 11, pp. 5450–5463, 2019.
- [54] R. Hadsell, S. Chopra, and Y. LeCun, “Dimensionality reduction by learning an invariant mapping,” in *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, vol. 2, 2006, pp. 1735–1742.
- [55] R. Chalapathy, A. K. Menon, and S. Chawla, “Anomaly detection using one-class neural networks,” *CoRR*, vol. abs/1802.06360, 2018. [Online]. Available: <http://arxiv.org/abs/1802.06360>
- [56] P. Oza and V. M. Patel, “One-class convolutional neural network,” *IEEE Signal Processing Letters*, vol. 26, no. 2, pp. 277–281, 2019.
- [57] G. Griffin, A. Holub, and P. Perona, “Caltech256 image dataset,” *Technical Report*, 2007. [Online]. Available: http://www.vision.caltech.edu/Image_Datasets/Caltech256/
- [58] Y. LeCun and C. Cortes, “MNIST handwritten digit database,” 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [59] P. Perera and V. M. Patel, “Learning deep features for one-class classification,” *CoRR*, vol. abs/1801.05365, 2018. [Online]. Available: <http://arxiv.org/abs/1801.05365>
- [60] M. Sabokrou, M. Khalooei, M. Fathy, and E. Adeli, “Adversarially learned one-class classifier for novelty detection,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3379–3388.
- [61] C. You, D. P. Robinson, and R. Vidal, “Provable self-representation based outlier detection in a union of subspaces,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 4323–4332.
- [62] P. Perera, R. Nallapati, and B. Xiang, “Ogan: One-class novelty detection using gans with constrained latent representations,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 2893–2901.
- [63] D. Abati, A. Porrello, S. Calderara, and R. Cucchiara, “Latent space autoregression for novelty detection,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 481–490.
- [64] M. C. Tsakiris and R. Vidal, “Dual principal component pursuit,” in *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*, 2015, pp. 850–858.
- [65] M. Kliger and S. Fleishman, “Novelty detection with gan,” 2018.
- [66] S. Fatemifar, S. Arashloo Rahimzadeh, M. Awais, and J. Kittler, “Client-specific anomaly detection for face presentation attack detection pattern recognition,” *Pattern Recognition*, 2020.
- [67] S. Rahimzadeh Arashloo, “Unseen face presentation attack detection using sparse multiple kernel fisher null-space,” *IEEE Transactions on Circuits and Systems for Video Technology*, 2020.
- [68] “Information technology — biometric presentation attack detection — part 3: Testing and reporting,” International Organization for Standardization, Standard, 2017.
- [69] Z. Boulkenafet, J. Komulainen, L. Li, X. Feng, and A. Hadid, “Oulu-npu: A mobile face presentation attack database with real-world variations,” in *2017 12th IEEE International Conference on Automatic Face Gesture Recognition (FG 2017)*, May 2017, pp. 612–618.
- [70] Y. Qin, W. Zhang, J. Shi, Z. Wang, and L. Yan, “One-class adaptation face anti-spoofing with loss function search,” *Neurocomputing*, vol. 417, pp. 384 – 395, 2020.
- [71] H. Feng, Z. Hong, H. Yue, Y. Chen, K. Wang, J. Han, J. Liu, and E. Ding, “Learning generalized spoof cues for face anti-spoofing,” *CoRR*, vol. abs/2005.03922, 2020.
- [72] H. Chen, G. Hu, Z. Lei, Y. Chen, N. M. Robertson, and S. Z. Li, “Attention-based two-stream convolutional networks for face spoofing detection,” *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 578–593, 2020.
- [73] W. Sun, Y. Song, C. Chen, J. Huang, and A. C. Kot, “Face spoofing detection based on local ternary label supervision in fully convolutional networks,” *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3181–3196, 2020.
- [74] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 2.1,” <http://cvxr.com/cvx>, Mar. 2014.
- [75] G.W. Stewart, *Matrix algorithms – Volume I: Basic decompositions*. SIAM, 2001.
- [76] S. R. Arashloo and J. Kittler, “Class-specific kernel fusion of multiple descriptors for face verification using multiscale binarised statistical image features,” *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2100–2109, Dec 2014.
- [77] P. Benner, P. Ezzatti, E. Quintana-Ortí, and A. Remón, “Matrix inversion on cpu-gpu platforms with applications in control theory,” *Concurrency and Computation: Practice and Experience*, vol. 25, no. 8, pp. 1170–1182, 2013. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.2933>
- [78] D. Yu, S. He, Y. Huang, G. Yu, and L. Yang, “A fast parallel matrix inversion algorithm based on heterogeneous multicore architectures,” in *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2015, pp. 903–907.

Supplementary Material

Shervin Rahimzadeh Arashloo

I.

In this section, we show that the projection $\phi(x_i)^\top \mathbf{w}$ defined in terms of the optimal solution of Eq. 5 yields a kernel Fisher null-space projection for one-class classification. The problem in Eq. 5 may be written as

$$\min_{\mathbf{w}} \frac{1}{n} \sum_{i=1}^n (1 - \phi(\mathbf{x}_i)^\top \mathbf{w})^2 = \min_{\mathbf{w}} \frac{1}{n} \|\mathbf{1} - \Phi(\mathbf{X})^\top \mathbf{w}\|_2^2$$

where $\mathbf{1}$ is an n -vector of 1's while $\Phi(\mathbf{X})$ denotes a matrix collection of positive training samples in the RKHS space. The optimal $\mathbf{w}$ may be found by taking the derivative of the cost function w.r.t $\mathbf{w}$ and setting it to zero, yielding

$$\mathbf{w} = (\Phi(\mathbf{X})\Phi(\mathbf{X})^\top)^{-1}\Phi(\mathbf{X})\mathbf{1}$$

Having found $\mathbf{w}$, the responses of the training observations (denoted as $\mathbf{y}$) may be determined as

$$\mathbf{y} = \Phi(\mathbf{X})^\top \mathbf{w} = \Phi(\mathbf{X})^\top (\Phi(\mathbf{X})\Phi(\mathbf{X})^\top)^{-1}\Phi(\mathbf{X})\mathbf{1} = \mathbf{B}\mathbf{1}$$

where $\mathbf{B} = \Phi(\mathbf{X})^\top (\Phi(\mathbf{X})\Phi(\mathbf{X})^\top)^{-1}\Phi(\mathbf{X})$. It can be easily shown that $\mathbf{B}$ is an identity matrix. For this purpose, let us compute $\Phi(\mathbf{X})\mathbf{B}$:

$$\Phi(\mathbf{X})\mathbf{B} = [\Phi(\mathbf{X})\Phi(\mathbf{X})^\top][\Phi(\mathbf{X})\Phi(\mathbf{X})^\top]^{-1}\Phi(\mathbf{X}) = \Phi(\mathbf{X})$$

Since $\Phi(\mathbf{X})\mathbf{B} = \Phi(\mathbf{X})$, hence $\Phi(\mathbf{X})(\mathbf{B} - \mathbf{I}) = \mathbf{0}$ and consequently $\Phi(\mathbf{X})^\top \Phi(\mathbf{X})(\mathbf{B} - \mathbf{I}) = \mathbf{0}$. However, $\Phi(\mathbf{X})^\top \Phi(\mathbf{X})$ is the kernel matrix by definition and as a result $\mathbf{K}(\mathbf{B} - \mathbf{I}) = \mathbf{0}$. Using an RBF kernel and assuming that not all training samples coincide, the kernel matrix $\mathbf{K}$ shall be positive definite and invertible, and hence, $(\mathbf{B} - \mathbf{I}) = \mathbf{K}^{-1}\mathbf{0} = \mathbf{0}$ and as a result $\mathbf{B}$ is an identity matrix. Consequently, $\mathbf{y} = \mathbf{B}\mathbf{1} = \mathbf{1}$. That is, the responses $\mathbf{y}$ derived using the optimal solution of Eq. 5 are all equal to 1. In the next section, it is shown that obtaining equal responses for all target training samples is equivalent to the Fisher null-space method for OCC.

II.

In a Fisher classifier, the criterion function is

$$\mathbf{J} = \frac{\varphi^\top \mathbf{S}_b \varphi}{\varphi^\top \mathbf{S}_w \varphi}$$

For the two-category case, the Fisher criterion may be equivalently written as the ratio of the between-class variance to the within-class variance [1]:

$$\mathbf{J} = \frac{(M_1 - M_2)^\top (M_1 - M_2)}{S_1^2 + S_2^2}$$

where M_1 and M_2 denote the means of the two classes while S_1 and S_2 represent the within-class variances. The analysis here employs the second equivalent representation of the Fisher criterion. As M_1 denotes the mean of the positive samples after projection (i.e. the mean of $\mathbf{y}$), we have $M_1 = 1$. As a result, the within-class variance of the projected positive samples is

$$S_1^2 = \sum_{i=1}^n (y_i - M_1)^\top (y_i - M_1) = \sum_{i=1}^n (1 - 1)^\top (1 - 1) = 0$$

where y_i denotes the projection of sample x_i onto an optimal feature subspace. As for the negative class, it is assumed that only a single hypothetical sample exists at the origin, thus $M_2 = 0$ and $S_2 = 0$.

Regarding the numerator of $\mathbf{J}$ one has

$$(M_1 - M_2)^\top (M_1 - M_2) = (1 - 0)^\top (1 - 0) = 1$$

Consequently, the between-class scatter is positive (i.e. $(M_1 - M_2)^\top (M_1 - M_2) = 1$) while the total within-class scatter is zero (i.e. $S_1^2 + S_2^2 = 0$) and hence the projection $\phi(\mathbf{x})^\top \mathbf{w}$ that minimises Eq. 5 yields a kernel Fisher null-space projection for OCC.

III.

In this section, we analyse the convergence properties of the proposed ℓ_p -norm multiple kernel Fisher null-space approach summarised as Algorithm 1. Let us first consider the following definition and the theorem that follows.

Definition If there exists a real number $L \geq 0$ such that $\forall x, x' \in X$, function $f : X \rightarrow Y$ satisfies

$$\|f(x) - f(x')\|_2 \leq L \|x - x'\|_2$$

then $f(\cdot)$ is Lipschitz continuous over X and L is a Lipschitz constant of $f(\cdot)$ [2]. The following theorem from [3] provides sufficient conditions for the convergence of a fixed-point iteration approach.

Theorem 1 If $f(\cdot)$ is continuous Lipschitz with $L < 1$ then:
i) the sequence $\{x_m, m = 1, 2, \dots\}$ generated by the iteration $x_{m+1} = f(x_m)$ will converge towards a fixed-point for any initial guess x_0 of the solution;
ii) the fixed point is unique.

Proof

i) As $f(\cdot)$ is Lipschitz continuous and $L < 1$, then for $x_m, m = 0, 1, 2, \dots$, we have

$$\begin{aligned} \|x_2 - x_1\|_2 &= \|f(x_1) - f(x_0)\|_2 \leq L \|x_1 - x_0\|_2 \\ \|x_3 - x_2\|_2 &= \|f(x_2) - f(x_1)\|_2 \leq L \|x_2 - x_1\|_2 \\ &\vdots \\ \|x_{m+1} - x_m\|_2 &= \\ &\|f(x_m) - f(x_{m-1})\|_2 \leq L \|x_m - x_{m-1}\|_2 \end{aligned}$$

Combining the inequalities above yields

$$\|x_{m+1} - x_m\|_2 \leq L^m \|x_1 - x_0\|_2$$

As $L < 1$ then $\lim_{m \rightarrow \infty} L^m = 0$, and consequently, $\{x_m, m = 1, 2, \dots\}$ converges to a point.

ii) Let us suppose that there are two distinct fixed-points x^* and $x^\bullet$. Since the function is continuous Lipschitz and $L < 1$, we have:

$$\begin{aligned} \|x^* - x^\bullet\|_2 &= \\ \|f(x^*) - f(x^\bullet)\|_2 &\leq L \|x^* - x^\bullet\|_2 < \|x^* - x^\bullet\|_2 \end{aligned}$$

Thus, we have arrived at the contradiction $\|x^* - x^\bullet\|_2 < \|x^* - x^\bullet\|_2$. As a result, there cannot be two distinct fixed points and the fixed point is unique.

We now provide the convergence requirements of the proposed method in the following proposition followed by its proof.

Theorem 2 For a sufficiently large δ , $f(\alpha) = \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j \mathbf{K}_j \right)^{-1} \mathbf{1}$ is continuous Lipschitz and its Lipschitz constant will be smaller than 1, and as a result, due to Theorem 1, the proposed approach will converge towards a fixed point which is unique irrespective of the initialisation of α .

Proof

For the proof, it suffices to show that for a sufficiently large δ , it holds that

$$\frac{\|f(\alpha^{m+1}) - f(\alpha^m)\|_2}{\|\alpha^{m+1} - \alpha^m\|_2} < 1$$

where α^m and α^{m+1} correspond to two consecutive solutions produced by the proposed algorithm for iteration m and iteration $m+1$. Let's suppose $\|\alpha^{m+1} - \alpha^m\|_2 = \omega$. If $\omega = 0$, it follows that the sequence α^m is converged. Hence, we consider the case when ω is strictly positive, i.e. $\omega > 0$. Considering the numerator of the quotient in the equation above, we have

$$\begin{aligned} \|f(\alpha^{m+1}) - f(\alpha^m)\|_2 &\leq \|f(\alpha^{m+1})\|_2 + \|f(\alpha^m)\|_2 \\ &= \left\| \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^{m+1} \mathbf{K}_j \right)^{-1} \mathbf{1} \right\|_2 + \left\| \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^m \mathbf{K}_j \right)^{-1} \mathbf{1} \right\|_2 \\ &\leq \sqrt{n} \left\| \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^{m+1} \mathbf{K}_j \right)^{-1} \right\|_2 + \sqrt{n} \left\| \left(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^m \mathbf{K}_j \right)^{-1} \right\|_2 \end{aligned}$$

where the first relation above is due to the triangle inequality and the last inequality follows from the fact that the norm of a matrix-vector product is upper bounded by the product of the norms of the matrix and the vector [4]. Moreover, it is known that the l_2 -norm of a square matrix is equal to the largest eigenvalue of the matrix and the l_2 -norm of the inverse of a matrix is $1/\sigma$ where σ denotes the minimum eigenvalue [4]. Consequently

$$\|f(\alpha^{m+1}) - f(\alpha^m)\|_2 \leq \sqrt{n}(1/\sigma_1 + 1/\sigma_2)$$

where σ_1 is the smallest eigenvalue of $(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^{m+1} \mathbf{K}_j)$ and σ_2 corresponds to the smallest eigenvalue of $(\delta \mathbf{I} + \sum_{j=1}^J \beta_j^m \mathbf{K}_j)$. Adding a constant to the elements on the main diagonal of a matrix shifts the eigenvalues by the same constant. For the proof, let us assume matrix $\mathbf{K}$ has an eigenvector $\mathbf{z}$ with the corresponding eigenvalue η , hence $\mathbf{K}\mathbf{z} = \eta\mathbf{z}$. If one adds a constant c to the main diagonal elements of $\mathbf{K}$, it yields $\mathbf{K}' = \mathbf{K} + c\mathbf{I}$ where $\mathbf{I}$ denotes an identity matrix. Let us calculate $\mathbf{K}'\mathbf{z}$:

$$\mathbf{K}'\mathbf{z} = (\mathbf{K} + c\mathbf{I})\mathbf{z} = \mathbf{K}\mathbf{z} + c\mathbf{I}\mathbf{z} = \eta\mathbf{z} + c\mathbf{z} = (\eta + c)\mathbf{z}$$

As a result, $\mathbf{z}$ is an eigenvector of $\mathbf{K}'$ with the eigenvalue of $\eta + c$.

Consequently, if the minimum eigenvalues of the combined kernel matrices $\sum_{j=1}^J \beta_j^{m+1} \mathbf{K}_j$ and $\sum_{j=1}^J \beta_j^m \mathbf{K}_j$ are denoted as σ'_1 and σ'_2 , respectively, we have $\sigma_1 = \sigma'_1 + \delta$ and $\sigma_2 = \sigma'_2 + \delta$. In the worst case scenario where the combined kernel matrices' minimum eigenvalues are zero, one has $\sigma_1 = \delta$ and $\sigma_2 = \delta$ and hence

$$\|f(\alpha^{m+1}) - f(\alpha^m)\|_2 \leq 2\sqrt{n}/\delta$$

Plugging the upper bound above into the quotient where start from gives the sufficient condition for the convergence of the proposed algorithm to a unique fixed point:

$$\frac{2\sqrt{n}/\delta}{w} < 1$$

or equivalently

$$\delta > \frac{2\sqrt{n}}{w}$$

For a sufficiently large δ , the above inequality holds true and as a result $f(\alpha)$ would be Lipschitz continuous with a Lipschitz constant smaller than one. Consequently, according to Theorem 1, the proposed algorithm would converge to a fixed point which is unique irrespective of the initialisation. ■

Using a Gaussian kernel and presuming that no duplicate training observations exist, the fused kernel matrix shall become strictly positive-definite. In this situation, the minimum bound derived for δ which will ensure convergence of the algorithm will be smaller than that given above since σ_1 and σ_2 will be larger than δ by σ'_1 and σ'_2 , respectively.

IV.

In this section, it is shown that $\beta^\top \mathbf{u} = \|\mathbf{u}\|_{p/(p-1)}$. From Eq. 18, we have $\beta = \frac{\mathbf{u}^{1/(p-1)}}{\|\mathbf{u}^{1/(p-1)}\|_p}$ and hence:

$$\begin{aligned} \beta^\top \mathbf{u} &= \frac{[\mathbf{u}^{1/(p-1)}]^\top \mathbf{u}}{\|\mathbf{u}^{1/(p-1)}\|_p} = \frac{\sum_{j=1}^J u_j^{1/(p-1)} u_j}{\left[\sum_{j=1}^J u_j^{p/(p-1)} \right]^{1/p}} \\ &= \frac{\left\{ \left[\sum_{j=1}^J u_j^{p/(p-1)} \right]^{(p-1)/p} \right\}^{p/(p-1)}}{\left\{ \left[\sum_{j=1}^J u_j^{p/(p-1)} \right]^{(p-1)/p} \right\}^{p/(p-1) \times 1/p}} \\ &= \frac{\|\mathbf{u}\|_{p/(p-1)}^{p/(p-1)}}{\|\mathbf{u}\|_{p/(p-1)}^{1/(p-1)}} = \|\mathbf{u}\|_{p/(p-1)} \end{aligned}$$

By plugging the above equality into the cost function, the objective function shall be

$$\max_{\alpha} -\delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} - \|\mathbf{u}\|_{p/(p-1)}$$

V.

The gradient of the cost function in Eq. 21 may be derived as

$$\begin{aligned} &\frac{\partial}{\partial \alpha} \left[-\delta \alpha^\top \alpha + 2\alpha^\top \mathbf{1} - \|\mathbf{u}\|_{p/(p-1)} \right] \\ &= -2\delta \alpha + 2 \times \mathbf{1} - \frac{\partial}{\partial \alpha} \left[\sum_{j=1}^J u_j^{p/(p-1)} \right]^{(p-1)/p} \\ &= -2\delta \alpha + 2 \times \mathbf{1} - (p-1)/p \times \left[\sum_{j=1}^J \frac{\partial}{\partial \alpha} (\alpha^\top \mathbf{K}_j \alpha)^{p/(p-1)} \right] \\ &\times \left[\sum_{j=1}^J u_j^{p/(p-1)} \right]^{-1/p} \\ &= 2 \left[-\delta \alpha + \mathbf{1} - \|\mathbf{u}\|_{p/(p-1)}^{-1/(p-1)} \left(\sum_{j=1}^J u_j^{1/(p-1)} \mathbf{K}_j \right) \alpha \right] \end{aligned}$$

A gradient ascent approach to perform optimisation shall be $\alpha_{t+1} = \alpha_t + \kappa \frac{\partial}{\partial \alpha} Cost$ where κ stands for the step size.

VI. EFFECT OF DIFFERENT CHOICES FOR p

The goal of the experiments in this section is to illustrate different behaviour of the proposed ℓ_p MK-FN approach with regards to different choices for p . To this end, two representative values of p , namely, $p = 2$ and $p = 1$ are considered that lead to non-sparse and sparse kernel weights, respectively.

A. Synthetic dataset

In the first part of the experiments in this section, we use a synthetic dataset. We form the target class by sampling 1000 observations from a 2D Gaussian distribution where the mean of the distribution in each dimension is randomly selected between 0 and 1 from a uniform distribution. The covariance

Fig. 1. Different characteristics of the ℓ_1 - and ℓ_2 -norm regularisation in the proposed method: mean and standard deviation of the classification error as a function of number kernels.

Fig. 2. Comparison of the performance of a sparse (ℓ_1 -norm) and a non-sparse (ℓ_2 -norm) regularisation scheme in the proposed approach in the presence of noisy kernels on the UCI Multiple Features dataset.

of the target distribution is randomly constructed. The test data is formed by sampling 1000 test points, 500 samples of which are drawn from the same distribution as that of the target class to form the positive test data and 500 samples from a different distribution whose mean is drawn from a uniform distribution between 0.5 and 1.5 to form the negative test observations. A validation set consisting of 100 positive and 100 negative samples is randomly generated. The positive validation samples are drawn from the same distribution as that of the target class while the negative validation samples are drawn from the same distribution as that of the negative test data. We repeat the random process above J times, yielding J kernels that can be assumed as kernels each capturing a distinct "view" of the same OCC problem. With this picture in mind, the ℓ_1 and ℓ_2 MK-FN approach are used to learn kernel weights. J is varied from 1 to 50 with a step of 5 where at each step the error rates corresponding to the ℓ_1 and ℓ_2 MK-FN methods are recorded. The process is above repeated 100 times for each choice of J . The mean and standard deviation of the classification errors are plotted in Figure 1. From Fig. 1, it can be observed that while increasing the number of kernels reduces the error rate for the ℓ_2 regularisation considerably, it does not have a big impact on the performance of the ℓ_1 regularisation. This is understandable as using more kernels is equivalent to bringing in more information and the ℓ_2 -norm may benefit from the additional information by assigning non-sparse weights to different kernels. The ℓ_1 -norm, on the other hand, does not benefit as much from such an increase in the number of kernels due to its over-selective behaviour.

TABLE I

COMPARISON OF THE PERFORMANCE OF THE PROPOSED APPROACH AGAINST TEN MULTIPLE KERNEL FUSION METHODS ON THE DELFT ONE-CLASS DATASETS IN TERMS OF AUC (MEAN $\pm$ STD %).

	Vehicle	Diabetes	Sonar	Breast	Imports	Glass	Wine	Liver	Survival
FN-Ave.	90.8 $\pm$ 3.9	57.6 $\pm$ 4.9	57.1 $\pm$ 3.9	97.1 $\pm$ 1.8	66.4 $\pm$ 8.3	83.3 $\pm$ 6.0	99.3 $\pm$ 0.3	62.6 $\pm$ 5.8	66.7 $\pm$ 7.0
FN-Prod.	90.8 $\pm$ 2.1	57.7 $\pm$ 6.5	60.5 $\pm$ 4.6	92.5 $\pm$ 4.0	73.8 $\pm$ 9.2	82.8 $\pm$ 6.3	99.2 $\pm$ 1.0	61.3 $\pm$ 5.9	62.6 $\pm$ 7.5
GP-Ave.	82.1 $\pm$ 3.1	58.8 $\pm$ 4.4	63.8 $\pm$ 5.4	84.6 $\pm$ 5.9	60.5 $\pm$ 7.1	82.9 $\pm$ 6.0	99.6 $\pm$ 0.4	61.1 $\pm$ 6.0	57.9 $\pm$ 5.4
GP-Prod.	39.8 $\pm$ 8.0	44.7 $\pm$ 5.7	47.0 $\pm$ 8.6	30.8 $\pm$ 11.4	55.4 $\pm$ 12.0	49.8 $\pm$ 10.9	65.0 $\pm$ 12.7	56.5 $\pm$ 6.7	47.2 $\pm$ 9.1
KPCA-Ave.	78.7 $\pm$ 4.1	50.2 $\pm$ 5.4	54.1 $\pm$ 4.6	96.6 $\pm$ 1.8	66.6 $\pm$ 8.2	74.8 $\pm$ 8.4	99.5 $\pm$ 0.3	50.9 $\pm$ 6.9	67.2 $\pm$ 7.2
KPCA-Prod.	82.2 $\pm$ 3.6	49.4 $\pm$ 6.4	57.4 $\pm$ 4.8	86.7 $\pm$ 4.9	75.1 $\pm$ 9.2	74.7 $\pm$ 8.2	99.1 $\pm$ 0.9	50.9 $\pm$ 6.9	64.9 $\pm$ 7.1
MK-SVDD	84.9 $\pm$ 5.2	51.5 $\pm$ 10.2	57.0 $\pm$ 11.2	80.9 $\pm$ 12.7	50.0 $\pm$ 0.0	75.4 $\pm$ 14.3	93.5 $\pm$ 6.1	57.6 $\pm$ 7.2	62.3 $\pm$ 5.9
Slim-MK-SVDD	89.8 $\pm$ 4.3	58.1 $\pm$ 5.6	63.0 $\pm$ 7.2	96.5 $\pm$ 1.6	72.4 $\pm$ 8.1	82.3 $\pm$ 7.0	99.2 $\pm$ 0.9	58.4 $\pm$ 7.3	65.7 $\pm$ 6.3
MK-OCSVM	84.9 $\pm$ 5.1	51.9 $\pm$ 9.8	56.6 $\pm$ 11.1	84.7 $\pm$ 8.7	53.2 $\pm$ 10.3	77.2 $\pm$ 13.5	93.5 $\pm$ 5.8	57.4 $\pm$ 7.3	62.8 $\pm$ 5.6
Slim-MK-OCSVM	90.6 $\pm$ 4.0	58.0 $\pm$ 4.6	62.6 $\pm$ 7.5	96.8 $\pm$ 1.3	72.4 $\pm$ 8.2	81.8 $\pm$ 7.0	99.1 $\pm$ 0.9	58.4 $\pm$ 7.3	65.3 $\pm$ 6.5
ℓ_p MK-FN	92.8 $\pm$ 2.8	61.3 $\pm$ 4.6	66.0 $\pm$ 5.2	97.6 $\pm$ 1.3	81.4 $\pm$ 8.5	84.4 $\pm$ 5.8	99.7 $\pm$ 0.3	64.5 $\pm$ 5.4	70.5 $\pm$ 4.9

B. UCI Multiple Features dataset

In this section, we study the impact of different choices for p , namely, $p = 1$ and $p = 2$ on the behaviour of the proposed ℓ_p MK-FN approach in the presence of noisy and uninformative kernels on the Multiple Features digit recognition dataset¹ from the UCI machine learning repository. The dataset incorporates 6 feature representations for 2000 handwritten digits. As it is a multi-class dataset, we convert it into a one-class dataset by considering a single class as the normal/target class and all the others as anomalies with respect to the target class. We then repeat this procedure for each class in the dataset, resulting in 10 one-class problems. Using the six representations provided, we build six informative kernels. Additionally, we build six noisy kernels using randomly generated data samples and mix them with the informative ones. We build a multiple kernel fusion system by combining six kernels where the number of noisy kernels within the composite kernel is gradually increased from zero to six. We record the area under the ROC curve (AUC) over all one-class problems (all digits) as a measure of the average performance. The procedure above is repeated 100 times where in each run we randomly select 60% of the positive observations as target training samples, 20% as positive validation data and 20% as positive test observations. For each one-class problem, the negative test and the negative validation samples are selected from all classes but the class of interest. The results corresponding to this experiment are depicted in Fig. 2. From the figure, one may observe that the non-sparse regularisation outperforms the sparse variant for all different numbers of noisy kernels. In particular, while the ℓ_2 regularisation provides almost stable performance over a wide range of number of noisy kernels, the ℓ_1 regularisation exhibits relatively inferior performance by locking onto a single, and possibly noisy kernel when the number of uninformative kernels is increased. This is particularly evident when the number of noisy kernels is increased towards 4 or 5 where the ℓ_1 -norm provides a poor performance close to random guess whereas the non-sparse ℓ_2 -norm provides relatively much better resilience against noisy information in the system.

¹Available at <https://archive.ics.uci.edu/ml/support/Multiple+Features>

VII. EXPERIMENT ON BENCHMARK ONE-CLASS DATASETS

In this section, the results of an evaluation of the proposed method on several benchmark one-class datasets from the repository of the Delft university for one-class classification [5] are provided. These are the *Vehicle*, *Diabetes*, *Sonar*, *Breast*, *Imports*, *Glass*, *Wine*, *Liver* and *Survival* datasets. Following earlier studies [6], each feature attribute of the provided feature vectors is used to construct a separate kernel. On each dataset, 60% of the data is randomly selected for training while 20% is used for validation and the rest for testing purposes. Each experiment is repeated ten times and the mean and the standard deviations of AUCs' (area under the ROC curves) are reported in Table I. From Table I, it may be observed that the FN-average approach (Fisher null-space method operating on the arithmetic mean kernel) provides very competitive performance to the SVM-based multiple kernel learning methods. In particular, the FN-average method typically performs better than the SVM-based multiple kernel learning methods (i.e. MK-SVDD and MK-OCSVM) while providing a close performance to the "slim" variants of these approaches. This is achieved despite using a fixed rule for combining multiple kernels in the FN-average method which emphasises the effectiveness of the Fisher null-space approach for one-class classification. Furthermore, it may be observed that the proposed ℓ_p MK-FN approach outperforms other multiple kernel fusion methods on all datasets which illustrates the efficacy of learning kernel combination weights for OCC and also highlights the superiority of the proposed OCC MKL approach compared to other fix rule or support vector MKL methods.

REFERENCES

- [1] E. Alpaydin, *Introduction to Machine Learning*, 3rd ed., ser. Adaptive Computation and Machine Learning. Cambridge, MA: MIT Press, 2014.
- [2] M. O. Searcoid, *Metric Spaces*. Berlin, New York: Springer-Verlag, ISBN 978-1-84628-369-7: Springer undergraduate mathematics series, 2006.
- [3] J. Hoffman and S. Frankel, *Numerical Methods for Engineers and Scientists*, 2nd ed. CRC Press, 2001.
- [4] G.W. Stewart, *Matrix algorithms – Volume I: Basic decompositions*. SIAM, 2001.
- [5] "Delft university one-class dataset repository," <http://homepage.tudelft.nl/n9d04/occ/index.html>, accessed: 2020-06-01.

- [6] Y. Han, Y. Yang, X. Li, Q. Liu, and Y. Ma, “Matrix-regularized multiple kernel learning via (r,p) norms,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 10, pp. 4997–5007, 2018.