

Generalized One-Class Learning Using Pairs of Complementary Classifiers

Cherian, Anoop; Wang, Jue

TR2021-076 July 12, 2021

Abstract

One-class learning is the classic problem of fitting a model to the data for which annotations are available only for a single class. In this paper, we explore novel objectives for one-class learning, which we collectively refer to as Generalized One-class Discriminative Subspaces (GODS). Our key idea is to learn a pair of complementary classifiers to flexibly bound the one-class data distribution, where the data belongs to the positive half-space of one of the classifiers in the complementary pair and to the negative half-space of the other. To avoid redundancy while allowing non-linearity in the classifier decision surfaces, we propose to design each classifier as an orthonormal frame and seek to learn these frames via jointly optimizing for two conflicting objectives, namely: i) to minimize the distance between the two frames, and ii) to maximize the margin between the frames and the data. The learned orthonormal frames will thus characterize a piecewise linear decision surface that allows for efficient inference, while our objectives seek to bound the data within a minimal volume that maximizes the decision margin, thereby robustly capturing the data distribution. We explore several variants of our formulation under different constraints on the constituent classifiers, including kernelized feature maps. We demonstrate the empirical benefits of our approach via experiments on data from several applications in computer vision, such as anomaly detection in video sequences, human poses, and human activities. We also explore the generality and effectiveness of GODS for non-vision tasks via experiments on several UCI datasets, demonstrating state-of-the-art results.

IEEE Transactions on Pattern Analysis and Machine Intelligence

Generalized One-Class Learning Using Pairs of Complementary Classifiers

Anoop Cherian*

Jue Wang*

Abstract—One-class learning is the classic problem of fitting a model to the data for which annotations are available only for a single class. In this paper, we explore novel objectives for one-class learning, which we collectively refer to as *Generalized One-class Discriminative Subspaces* (GODS). Our key idea is to learn a pair of complementary classifiers to flexibly bound the one-class data distribution, where the data belongs to the positive half-space of one of the classifiers in the complementary pair and to the negative half-space of the other. To avoid redundancy while allowing non-linearity in the classifier decision surfaces, we propose to design each classifier as an orthonormal frame and seek to learn these frames via jointly optimizing for two conflicting objectives, namely: i) to minimize the distance between the two frames, and ii) to maximize the margin between the frames and the data. The learned orthonormal frames will thus characterize a piecewise linear decision surface that allows for efficient inference, while our objectives seek to bound the data within a minimal volume that maximizes the decision margin, thereby robustly capturing the data distribution. We explore several variants of our formulation under different constraints on the constituent classifiers, including kernelized feature maps. We demonstrate the empirical benefits of our approach via experiments on data from several applications in computer vision, such as anomaly detection in video sequences, human poses, and human activities. We also explore the generality and effectiveness of GODS for non-vision tasks via experiments on several UCI datasets, demonstrating state-of-the-art results.

Index Terms—one-class classification, subspace learning, kernelized subspaces, Riemannian optimization

1 INTRODUCTION

THERE are several real-world applications for which it may be straightforward to characterize the normal operating behavior and collect data to train learning systems, however may be difficult or sometimes even impossible to have data when abnormalities or rare events happen. Examples include but not limited to, an air conditioner making a spurious vibration, a network attacked by an intruder, sudden variations in a patient’s vitals, or an accident captured in a video surveillance camera, among others [1]. In machine learning literature, such problems are usually called one-class problems [2], [3], signifying the fact that we may have an unlimited supply of labeled training data for the one-class (corresponding to the normal operation of the system), but do not have any labels or training data for situations corresponding to abnormalities.

Popular problems in computer vision, such as video novelty detection [4], [5], [6], surveillance anomaly detection [7], [8], [9], image denoising [10], and outlier detection [11], [12], are variants of the standard one-class setting. The main goal of such methods is usually to learn a model that fits to the normal set, such that abnormalities can be characterized as outliers of this learned model. Given that the distribution of abnormal samples are often unbounded or could even partially overlap with the normal data samples (e.g., in a video surveillance application, when both

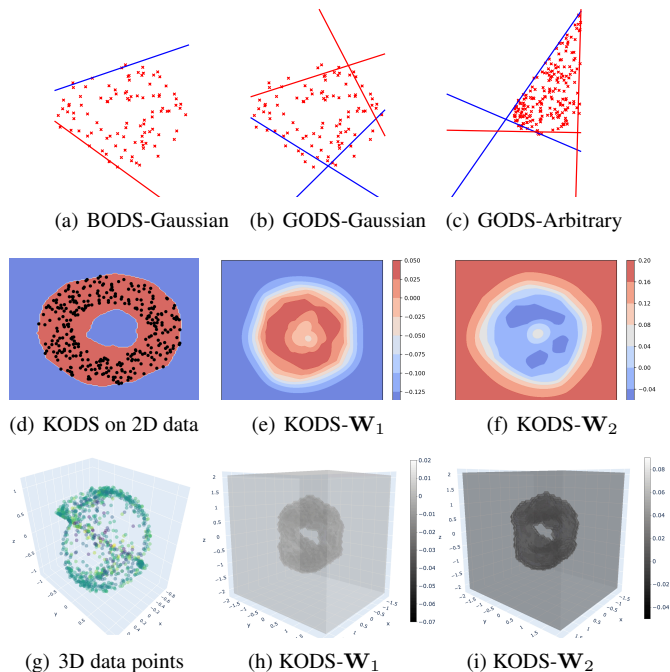


Fig. 1: Visualizations of decision regions using various GODS formulations on synthetic data. Figures (a, b, c) show subspaces found by BODS and GODS on various data distributions (see Section 7.5). The colors identify hyperplanes within a classifier in the complementary pair. Figure (d) shows KODS decision regions for ring-shaped data (black dots) using an RBF kernel. Figures (e, f) are the decision regions of the classifiers; \mathbf{W}_1 bounding data from outside and \mathbf{W}_2 from inside, together they define the region in (d). Figure (g,h,i) show 3D points and the decision surfaces of the two classifiers.

• Anoop Cherian (corresponding author) is with Mitsubishi Electric Research Labs (MERL), Cambridge, MA, E-mail: cherian@merl.com

• Jue Wang is with the Research School of Engineering, The Australian National University, ACT 2601, Australia. E-mail: jue.wang@anu.edu.au. Work done while interning at MERL.

* Equal contribution.

normal and abnormal features could be present in a time window), achieving one-class classification can be practically challenging.

Classical solutions to one-class problems are mainly extensions to support vector machines (SVMs), such as the one-class SVM (OC-SVM), that maximizes the margin of the discriminative hyperplane from the origin [13]. There are extensions of this scheme, such as the least-squares one-class SVM (LS-OSVM) [14] or its online variants [15], that learn to find a tube of minimal diameter that includes all the labeled data. Another popular approach is the support-vector data description (SVDD) that finds a hypersphere of minimum radius that encapsulates the training data [16]. There have also been kernelized extensions of these schemes that use the kernel trick to embed the data points in a reproducible kernel Hilbert space, potentially enclosing the ‘normal’ data with arbitrarily-shaped boundaries.

Apart from the classic one-class solutions, there is an increasing number of recent works that use deep neural networks for building the one-class model [6], [9], [17], [18], [19], [20]. In these approaches, typically a deep auto-encoder model is trained on the one-class data such that its reconstruction error is minimized. When such a model is provided with an out-of-distribution data sample (anomaly), the reconstruction error can be large, which could be used as an anomaly cue [6], [20]. There are extensions of this general architecture using generative adversarial networks (GANs) to characterize the distribution of the in-class samples [21], [22]. For anomaly detection on time-varying inputs, there are also adaptations using predictive auto-encoders, such as [9], that attempts to generate the (latent) future samples, and flags anomalies if the predicted sample is significantly different from the observed one.

While, these approaches have been widely adopted in several applications (see e.g., Chandala et al. [1]), they have drawbacks. For example, the OC-SVM uses only a single hyperplane, however using multiple hyperplanes may be beneficial [23]. The SVDD scheme makes a strong assumption on the spherical nature of the data distribution. Using kernel methods may impact scalability, while deep learning methods may need specialized hardware (such as GPUs) and large training sets. Thus, trading-off between the pros and cons of these diverse prior methods, we propose novel generalizations of these techniques, which we call *generalized one-class discriminative subspaces* (GODS). The key goal of GODS is to combine the linearity properties of OC-SVM, and the non-linear bounded characterization of SVDD in a single framework. However, our proposed one-class model is neither linear nor uses a spherical classifier, instead uses a pair of orthonormal frames¹ whose columns characterize linear classifiers. Specifically, these columns are optimized such that the one-class data belongs to the positive half-spaces of the columns in one of these classifiers and to the negative half-spaces of the columns in the other; thus these classifiers jointly form a complementary pair. These classifiers, with their respective half-spaces defined by their orthonormal columns, non-linearly bound the data from different directions. Our learning objective, to find these complementary classifiers, jointly optimizes two opposing criteria: i) to minimize the distance between the two classifiers, thus bounding the data within the smallest volume, and ii) to maximize the margin between the hyperplanes and the data, thereby avoiding overfitting, while improving classification robustness.

1. Orthonormal frames are matrices with linearly independent unit norm columns.

Our proposed GODS model offers several advantages against prior methods: (i) the piecewise linear decision boundaries approximate a non-linear classifier, while providing computationally cheap inference, and (ii) the use of the complementary classifier pair allows flexible bounding of the data distribution of arbitrary shapes, as illustrated in Figure 1. For example, our kernelized variant of GODS (called KODS), that we introduce in Section 4.4, bounds data from outside as well as inside (see Figure 1(d)), which is usually not possible in prior methods. Albeit these benefits, our objective is non-convex due to the orthogonality constraints. However, such non-convexity fortunately is not a significant practical concern as they naturally place the optimization objective on the Stiefel manifold [24]. This is a well-studied Riemannian manifold [25] for which there exist efficient non-linear optimization methods at our disposal. We use one such optimization scheme, called Riemannian conjugate gradient [26], which is fast and efficient.

To empirically evaluate the benefits of our GODS formulations, we apply them to one-class data arising from various anomaly detection problems in computer vision and machine learning. One novel task we consider in this paper is that of out-of-pose (OOP) detection in cars [27], [28]. Specifically, in this task, our goal is to detect if the passengers or the driver are seated OOP as captured by an inward-looking dashboard camera. For this task, we showcase the effectiveness of our approaches on a new dataset, which we call Dash-Cam-Pose. Apart from this task, we also report experimental results on several standard and public anomaly detection benchmarks in computer vision, such as on the UCF-crime [29] and the UCSD Ped2 [30] datasets. We also provide experiments on the standard JHMDB action recognition dataset [31] re-purposed for the anomaly detection task. We further analyze the generalizability of our approach to non-computer vision applications by providing results on five UCI datasets. Our results demonstrate that GODS variants lead to significant performance improvements over the state of the art.

We summarize below the main contributions of this paper:

- We first introduce a basic one-class discriminative subspace (BODS) classifier that uses a pair of hyperplanes.
- We generalize BODS to use multiple hyperplanes, termed generalized one-class discriminative subspaces (GODS) and derive a kernelized variant, termed KODS. We also present several formulations of GODS under different assumptions on the classifiers.
- We explore Riemannian conjugate gradient algorithms for optimizing our objectives. Specifically, the BODS and GODS formulations use a Stiefel manifold, while KODS is modeled on the generalized Stiefel manifold.
- We present a novel task of out-of-pose detection, and provide a new dataset, termed Dash-Cam-Pose.
- We provide experiments on Dash-Cam-Pose, three standard vision datasets, and five UCI datasets, demonstrating substantial performance benefits of our methods.

We note that this work is an extension of the ICCV conference paper [32] and extends it in the following ways: (i) we provide novel extensions to the GODS using different assumptions on the complementary classifiers (Section 4.3), (ii) we derive a kernelized variant of GODS (Section 4.4) and provide practical simplifications for optimization on the generalized Stiefel manifold (Section 6.2), (iii) we provide adaptive classification rules in Section 5, and (iv) provide additional experiments and ablation studies, including new results on UCSD Ped2 and UCI datasets.

2 RELATED WORK

As one-class problems arise in numerous practical settings, they have been explored to great depth in a variety of disciplines, including but not limited to remote sensing [33], network intruder detection [34], and fraud detection [35]. In computer vision, a few illustrative problems are novelty detection [4], [5], [6], video anomaly detection [7], [8], [36], diagnosis on medical images [37], and anomalous object attribute recognition in image collections [38]. For a comprehensive review of applications, we refer the interested reader to surveys, such as [1], [39], [40].

Classic methods for modelling such one-class problems are extensions of data density estimation techniques [41], [42]. These methods attempt to model the density of the given data in the input space by trading-off between maximizing their inclusivity within a (given) density quantile while minimizing its volume. The dependence on minimal density volume in the input space is discarded in Scholkopf et al. [13], [16] against smoothness of the decision function in a non-linear (kernelized) feature space. Working in the kernel space not only allows for more flexible characterizations of the distribution of the data samples, but also allows transferring the max-margin machinery (and the associated theory) developed for support vector machines to be directly used in the one-class setting. However, as alluded to earlier, working with kernel feature maps can be demanding in large data settings, and thus our main focus in this paper is on deriving one-class algorithms in the input (primal) space. That said, we also explore a kernelized dual variant of our scheme for problems that are impossible to be modelled using our primal variant.

Modern efforts to one-class learning typically use either (i) good hand-crafted representations combined with effective statistical learning models, or (ii) implicit representation and learning via neural networks. Below, we review these efforts in detail.

Explicit Modelling Approaches. Performance of any one-class approach inevitably depends on the effectiveness of the data representation. For example, visual representations such as histogram of oriented gradients [43] (HOG) and histogram of optical flows (HOF) [44] have been beneficial in developing several anomaly detection algorithms. A Markov random field (MRF) on HOG and HOF descriptors is proposed in Zhang et al. [45] for modeling the normal patterns in a semi-supervised manner, where an abnormal sample model is iteratively derived from the normal one using Bayesian adaptation. In Xu and Caramanis [46], an outlier pursuit algorithm is proposed using convex optimization for the robust PCA problem. Similarly, Kim et al. [47] propose a space-time MRF to detect abnormal activities in videos. This method uses a mixture of probabilistic principal component analysis to characterize the distribution of normal data characterized as densities on optical flow. Detecting out-of-context objects is explored in [48], [49] using support graph and generative models. Motion trajectory analysis [50], [51], [52] of objects in video sequences has been a common approach for modeling anomalies, under the strong assumption that deviant trajectories may correspond to abnormal data. Detecting salient regions in images has also been implemented by some researchers [53], [54]. In contrast to these approaches that propose problem-specific anomaly detection models, our solution is for a general setting.

Sparse reconstruction analysis has been a powerful workhorse in the recent times in developing several one-class solutions [55], [56], [57], [58]. The assumption in these methods is that normal data can be encoded as sparse linear combinations of columns in a dictionary that is learned only on the normal data; however the

reconstruction error of any out-of-distribution sample using this dictionary could be significant. In addition to the reconstruction loss, a study from Ren et al. [59] points out that the sparsity term should be taken into consideration for improving the anomaly detection accuracy. However, sparse reconstruction methods can be computationally expensive. To improve their efficiency, Bin et al. [57] proposes an online detection scheme using sparse reconstructibility of query signals from an atomically learned event dictionary. Yang et al. [55] improves efficiency via learning multiple small dictionaries to encode image patterns at multiple scales. While, our proposed GODS algorithm could also be treated as a dictionary of orthonormal columns, our inference is significantly cheaper in contrast to solving ℓ_1 -regularized problems in these works as GODS involves only evaluations of inner products of the data samples to the learned hyperplanes.

Deep Learning Based Methods. The huge success of deep neural networks on several fundamental problems in computer vision [60], [61] has also casted its impact in devising schemes for anomaly detection [20]. Extending classical methods, a deep variant of SVDD is proposed in Ruff et al. [62], however assumes the one-class data is unimodal. Variants of OC-SVM are explored in [18], [63]. Parera and Patel [64] proposes a trade-off between compactness and descriptiveness using an external reference dataset to train a deep model on one-class data. Liang et al. [65] proposes to use statistical trends in the softmax predictions. Deep learning based feature representations have been used as replacements for hand-crafted features in several one-class problem settings. For example, Xu et al. [66] design a multi-layer auto-encoder embracing data-driven feature learning. Similarly, Hasan et al. [20] propose a 3D convolutional auto-encoder to capture both spatial and temporal cues in video anomaly detection. Leveraging the success of convolutional neural networks (CNNs) to capture spatial cues, [67], [68], and [19] propose to embed recurrent networks, such as LSTMs, to model the appearance dynamics of normal data. In [69] and [70], frameworks to minimize the in-distribution sample distances is proposed thereby maximizing the distance to out-of-distribution samples. In Sabokrou et al. [71], a pre-trained CNN is used for extracting region features, and a cascaded outlier detection scheme is applied. Multiple instance learning (MIL) in a deep learning setting is attempted in [29] for anomaly detection using weakly-labeled training videos via applying an MIL ranking loss with sparsity and smoothness constraints; however includes both normal and abnormal samples in the training set. Deep generative adversarial networks (GAN) have also been proposed to characterize the single class [6], [21], [22], [72]. These methods typically follow the same philosophy of training auto-encoders, however uses an adversarial discriminator to improve quality of the decoded data sample; the discriminator confidence is then used as an abnormality cue during inference.

In contrast to these approaches, we focus on explicit modelling of one-class data distributions, allowing better and more controlled characterization of the single class. Deep learning approaches reviewed above are complimentary to our contributions; in fact we use deep-learned data representations in our experiments and simultaneously demonstrate our performances on non-deep-learned features as well. More recently, explicit characterization of the diversity of the normal data is explored in [73], self-supervised learning is proposed in [74], and normalizing flows for the task in [75]. We note that using our formulations within a deep neural network, while straightforward using methods described in [76], [77], is currently beyond the focus of this paper.

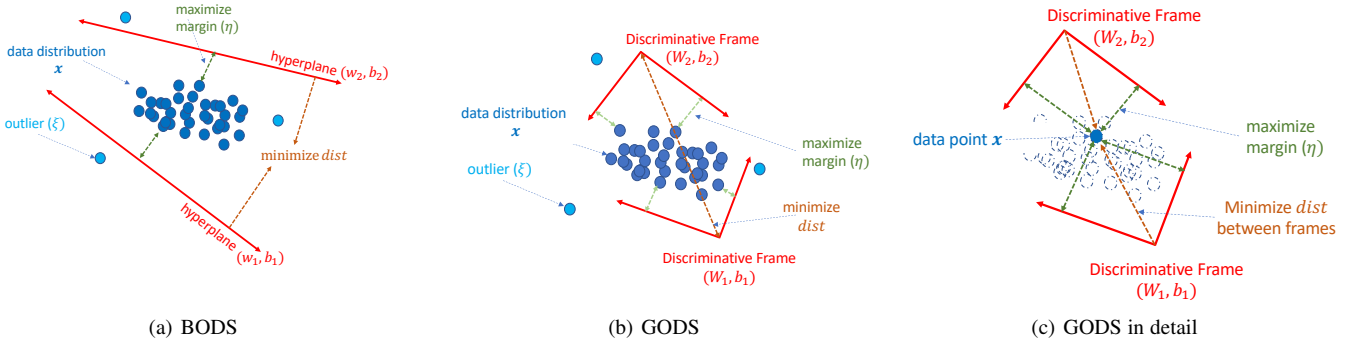


Fig. 2: (a), (b) graphically illustrate our BODS and GODS formulations. Our one-class variants learn discriminative hyperplanes (BODS) or orthonormal frames (GODS) that maximize the margin with the data distribution while minimizing the volume of the one-class region captured (via minimizing the distance dist^2 between hyperplanes or frames). In Figure 2(c), we show a detailed depiction of the objectives in GODS using a single data point \mathbf{x} . See Section 4.2 for more details.

3 BACKGROUND

Let $\mathcal{D} \subset \mathbb{R}^d$ denote the data distribution consisting of our one class-of-interest and everything outside it, denoted $\bar{\mathcal{D}}$, be the anomaly set. Suppose we are given n data instances $\mathcal{D}_o = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\} \subset \mathcal{D}$. The goal of one-class classifiers is to use \mathcal{D}_o to learn a functional f which is positive on \mathcal{D} and negative on $\bar{\mathcal{D}}$. Typically, the label of \mathcal{D} is assumed $+1$ and that of $\bar{\mathcal{D}}$ is -1 . We emphasize that we assume to have access *only* to the one-class data for training our models; i.e., the set $\bar{\mathcal{D}}$ is not available for training, and samples from $\bar{\mathcal{D}}$ are used only at test time.

In One-Class Support Vector Machine (OC-SVM) [13], f is modeled as an extension of an SVM objective by learning a max-margin hyperplane that separates the origin from the data points in \mathcal{D}_o . Mathematically, f has the form $\text{sgn}(\mathbf{w}^T \mathbf{x} + b)$, where $(\mathbf{w}, b) \in \mathbb{R}^d \times \mathbb{R}^1$ and is learned by minimizing the following objective:

$$\min_{\mathbf{w}, b, \xi \geq 0} \frac{1}{2} \|\mathbf{w}\|_2^2 - b + C \sum_{i=1}^n \xi_i, \text{ s.t. } \mathbf{w}^T \mathbf{x}_i + b + \xi_i \geq 0, \forall \mathbf{x}_i \in \mathcal{D}_o,$$

where ξ_i 's are non-negative slacks, b is the hyperplane intercept, and C is the slack penalty. As a single hyperplane in the input space might be insufficient to capture non-linear data, kernel feature maps are proposed in [13].

Another popular variant of one-class classifiers is the support vector data description (SVDD) [16] that instead of modeling data to belong to an open half-space of \mathbb{R}^d (as in the primal OC-SVM), assumes the labeled data inhabits a bounded set; specifically, the optimization seeks the centroid $\mathbf{c} \in \mathbb{R}^d$ of a hypersphere of minimum radius $R > 0$ that can contain all points (or a given quantile of points) in \mathcal{D}_o . Mathematically, the objective reads:

$$\min_{\mathbf{c}, R, \xi \geq 0} \frac{1}{2} R^2 + C \sum_{i=1}^n \xi_i, \text{ s.t. } \|\mathbf{x}_i - \mathbf{c}\|_2^2 \leq R^2 - \xi_i, \forall \mathbf{x}_i \in \mathcal{D}_o,$$

where, as in OC-SVM, the ξ 's model the slack. As is discussed in [13], OC-SVM and SVDD are equivalent when the underlying kernel is translation invariant. There have been extensions of this scheme, such as the mSVDD that uses a mixture of hyperspheres [78]. Other variants include (i) the density-induced SVDD [79], (ii) kernelized variants [80], and (iii) more recently, those that use subspaces for data description [81]. A major drawback of SVDD in general is the strong assumption it makes

on the isotropic nature of the underlying data distribution. Such a demand is ameliorated by combining OC-SVM with the idea of SVDD in least-squares one-class SVM (LS-OSVM) [14] that learns a tube around the discriminative hyperplane that contains the input; however, this scheme also makes strong assumptions on the data distribution (such as being cylindrical).

Unlike OC-SVM that learns a compact data model to enclose as many training samples as possible, a different approach is to use principal component analysis (PCA) (or its variants, such as Robust PCA [82], [83], [84], [85], [86], and Kernel PCA) to summarize the data by using its principal subspaces. However, such an approach is usually unfavorable due to its high computational cost, especially when the dataset is large. Similar in motivation to the proposed technique, Bodesheim et al. [87] uses a null space transform for novelty detection, while Liu et al. [88] optimizes a kernel-based max-margin objective for outlier removal and soft label assignment. However, their problem setups are different from ours in that [87] requires multi-class labels in the training data and [88] is proposed for unsupervised learning.

In contrast to these prior methods, we explore the one-class learning problem from a distinct perspective; specifically, to use orthonormal frames as in PCA, however instead of approximating the one-class data using these frames, our objective seeks a pair of complementary orthonormal frames that bounds the data discriminatively and in a piece-wise linear manner, such that the data subspace is sandwiched between these frames. We first present a simplified variant of this idea using two complementary hyperplanes, dubbed Basic One-class Discriminative Subspaces (BODS) (Fig. 2(a)); these hyperplanes are independently parameterized and bounds the data distribution. Note that there is a similar prior work, termed Slab-SVM [89], that learns two hyperplanes for one-class classification. However, their hyperplanes are constrained to have the same slope, which we do not impose in our BODS model; as a result, our model is more flexible than Slab-SVM. We extend BODS using multiple orthogonal hyperplanes, which we call Generalized One-class Discriminative Subspaces (GODS) (see Fig. 2(b)). The use of such discriminative subspaces has been recently explored in the context of representation learning on videos in Wang and Cherian [23] and Wang et al. [90], however requires a surrogate negative bag of features, that is found via adversarial learning.

Notation: We use bold-face upper-case for matrices, bold-face lower-case for vectors, non-bold-face lower case for scalars. We

use \mathbf{W} to represent an orthonormal frame. We sometimes succinctly refer to a set of variables without their subscripts; e.g., $\mathbf{W} = \{\mathbf{W}_1, \mathbf{W}_2\}$. The notation $\mathbf{1}_{p \times q}$ represents a matrix of ones with p rows and q columns.

4 PROPOSED METHOD

In this section, we formally introduce our schemes. First, we present BODS using a pair of hyperplanes, which we generalize to GODS using a pair of discriminative frames in Section 4.2. We explore variants of GODS in Section 4.3 and further generalize GODS using kernel feature maps, proposing KODS in Section 4.4. With a slight abuse of terminology, we call our entire suite of formulations as GODS.

4.1 Basic One-class Discriminative Subspaces

In this section, we introduce a basic variant of our objective, which we call Basic One-class Discriminative Subspaces (BODS). The key goal of BODS is to bound the one-class data distribution using a pair of hyperplanes. Similar to OC-SVM, we seek these hyperplanes to have a large margin from the data distribution, thus allowing robustness to minor differences in the test data distribution. Further, inspired by SVDD, we also demand the two hyperplanes to bound the data within a minimal volume. BODS combines these two conflicting objectives into a joint formulation. Mathematically, suppose (\mathbf{w}_1, b_1) and (\mathbf{w}_2, b_2) define the parameters of the pair of hyperplanes. Our goal in BODS is then to minimize an objective such that all data points \mathbf{x}_i be classified to the positive half-space of (\mathbf{w}_1, b_1) and to the negative half-space of (\mathbf{w}_2, b_2) , while also minimizing a suitable distance between the two hyperplanes (see Figure 2(a)). To this end, we propose the following BODS objective:

$$\min_{\substack{(\mathbf{w}_1, b_1), (\mathbf{w}_2, b_2), \\ \xi_1, \xi_2 \geq 0}} \frac{1}{2} \|\mathbf{w}_1\|_2^2 + \frac{1}{2} \|\mathbf{w}_2\|_2^2 - b_1 - b_2 + \mathbf{\Omega}(\xi_{1i}, \xi_{2i}) \\ + \frac{1}{2} \text{dist}^2((\mathbf{w}_1, b_1), (\mathbf{w}_2, b_2)), \quad (1)$$

$$\text{s.t. } \left(\mathbf{w}_1^T \mathbf{x}_i + b_1 \right) \geq \eta - \xi_{1i}, \quad (2)$$

$$\left(\mathbf{w}_2^T \mathbf{x}_i + b_2 \right) \leq -\eta + \xi_{2i}, \quad (3)$$

where the first four terms of (1) seek large margins similar to a standard OC-SVM objective. The key element of BODS is the last term in (1) that aims to minimize a suitable distance dist between the two hyperplanes. In (2) and (3), we capture the complementary classification constraints noted above. We use the notation $\mathbf{\Omega}(\xi_{1i}, \xi_{2i}) = C \sum_{i=1}^n (\xi_{1i}^2 + \xi_{2i}^2)$ for the squared-slack regularization² and $\eta > 0$ specifies a (given) classification margin. For BODS, we assume dist is the Euclidean distance, i.e., $\text{dist}^2((\mathbf{w}_1, b_1), (\mathbf{w}_2, b_2)) = \|\mathbf{w}_1 - \mathbf{w}_2\|_2^2 + (b_1 - b_2)^2$.

4.1.1 BODS on the Unit Sphere

It is often seen (especially for computer vision problems) that feature normalization, specifically unit normalizing of the inputs, demonstrate better performances [93], [94], [95]. While, such a step may be counter-intuitive as we lose the discriminativeness

in the scale of the input features, such normalization often allows counteracting the effects of *burstiness* [93] – a statistical phenomenon due to a non-uniform distribution of data elements. Such normalization is even found to improve the performances of deep-learned features when used in max-margin frameworks [96], [97] and its importance to one-class tasks is ascertained in [98], in the context of SVDD. As the primary focus of this paper is in developing one-class solutions for vision problems, we decided to make this normalization as part of the pre-processing steps when generating our input features and thus in deriving our core formulations, (albeit we consider variants of our formulation without such assumptions in the next section). We empirically validate this assumption in Section 7.5.

Following this idea, we assume that our data is unit normalized, i.e., $\|\mathbf{x}_i\|_2 = 1$, and thus belongs to a unit hypersphere U^{d-1} , which is a sub-manifold of the Euclidean manifold \mathbb{R}^d . This assumption³ on the data naturally places our hyperplanes also to belong to U^{d-1} ; i.e., $\|\mathbf{w}_1\|_2 = \|\mathbf{w}_2\|_2 = 1$. Using these manifold constraints, our BODS formulation can be rewritten as follows:

$$P_1 := \min_{\mathbf{w}_1, \mathbf{w}_2 \in U^{d-1}, b_1, b_2} \frac{1}{2} \alpha(b_1, b_2) - \mathbf{w}_1^T \mathbf{w}_2 \\ + \frac{\nu}{2n} \sum_i \left[\eta - \left(\mathbf{w}_1^T \mathbf{x}_i + b_1 \right) \right]_+^2 + \left[\eta + \left(\mathbf{w}_2^T \mathbf{x}_i + b_2 \right) \right]_+^2, \quad (4)$$

where $\alpha(b_1, b_2) = (b_1 - b_2)^2 - 2(b_1 - b_2)$. We further simplify the BODS objective by substituting the constraints on the slacks ξ 's in (2) and (3) into $\mathbf{\Omega}$ in (1) and include them in the objective as soft constraints using the hinge loss $[\]_+$. We use ν to denote a penalty factor on these soft constraints. In Fig. 2(a), we illustrate the decision boundaries of the BODS model.

While, BODS offers a simple and flexible model to capture the one-class distribution, the linearity of the classifiers may limit its applications to sophisticated data models. A natural idea is then to empower these classifiers to have non-linear decision boundaries. While, using a kernel method is perhaps a standard approach in this regard (which we present subsequently), we first propose to achieve non-linearity via piecewise linear decision boundaries. To this end, we equip each classifier in BODS with a set of hyperplanes; each set forming a complementary pair with the other. The use of piecewise linear decision boundaries makes inference computationally cheap as it requires only $2K$ inner products during inference, assuming K hyperplanes per set. Further, we also avoid the need for computing kernel matrices, allowing for scalability of our approach to larger datasets. However, using sets of hyperplanes brings in the challenge of how to effectively regularize them to avoid overfitting and redundancy. To this end, in the following subsections, we generalize BODS to use pairs of multiple hyperplanes, regularized as orthonormal frames, thus providing a richer and non-linear discriminative setup, and subsequently present other regularizations and kernel embeddings.

4.2 One-class Discriminative Subspaces

Let us continue to assume the input data is unit normalized, we will remove this assumption in the next section. Formally, suppose $\mathbf{W}_1, \mathbf{W}_2 \in \mathcal{S}_d^K$ be orthonormal frames – that is, matrices of dimensions $d \times K$, each with K columns where each column is

2. We use squared- ℓ_2 norm on the slacks instead of the standard ℓ_1 norm to allow for smooth gradients in our optimization setup. While, this choice may hurt the sparsity of the slacks, it is usually not empirically seen to affect the classifier generalization [91], [92].

3. Such an SVM formulation is classically known as *normalized margin SVMs*, theoretically analyzed for its generalization performance in [99] and [100][Sec. 10.6.2].

orthonormal to the rest; i.e., $\mathbf{W}_1^T \mathbf{W}_1 = \mathbf{W}_2^T \mathbf{W}_2 = \mathbf{I}_K$, where \mathbf{I}_K is the $K \times K$ identity matrix (see Fig. 2(b)). Such frames belong to the so-called Stiefel manifold, denoted \mathcal{S}_d^K , with K d -dimensional directions. Note that the orthogonality assumption on the \mathbf{W}_i 's is to ensure they capture diverse discriminative directions, leading to better regularization, while also improving their characterization of the data distribution. A direct extension of P_1 then leads to:

$$P_2 := \min_{\mathbf{W} \in \mathcal{S}_d^K, \mathbf{b}} \frac{1}{2} \text{dist}_{\mathbf{W}}^2(\mathbf{W}_1, \mathbf{W}_2) + \alpha(\mathbf{b}_1, \mathbf{b}_2) \quad (5)$$

$$+ \frac{\nu}{2n} \sum_i \left[\eta - \min(\mathbf{W}_1^T \mathbf{x}_i + \mathbf{b}_1) \right]_+^2 \quad (6)$$

$$+ \frac{\nu}{2n} \sum_i \left[\eta + \max(\mathbf{W}_2^T \mathbf{x}_i + \mathbf{b}_2) \right]_+^2, \quad (7)$$

where $\text{dist}_{\mathbf{W}}$ is a suitable distance between orthonormal frames, and $\mathbf{b} \in \mathbb{R}^K$ is a vector of biases, one for each hyperplane. Note that in (6) and (7), unlike BODS, $\mathbf{W}^T \mathbf{x}_i + \mathbf{b}$ is a K -dimensional vector. Thus, (6) says that the minimum value of this vector should be greater than η and (7) says that the maximum value of it is less than $-\eta$. To simplify the notation, let us use $\zeta(\mathbf{W}, \mathbf{b}) = \alpha(\mathbf{b}_1, \mathbf{b}_2) + (6) + (7)$. Then, P_2 can be written as follows:

$$P'_2 := \min_{\mathbf{W} \in \mathcal{S}_d^K, \mathbf{b}} -\text{Tr}(\mathbf{W}_1^T \mathbf{W}_2) + \zeta(\mathbf{W}, \mathbf{b}). \quad (8)$$

The formulation P'_2 , due to the first term, enforces a *tight coupling* between \mathbf{W}_1 and \mathbf{W}_2 ; such a coupling might prevent the frames from freely aligning to the data distribution, resulting in sub-optimal performance. To circumvent this issue, we propose the following work around. Recall that the main motivation to define the distance between the frames is so that they sandwich the (one-class) data points compactly. Thus, rather than defining a distance between the frames directly, one could also use a measure that minimizes the Euclidean distance of each data point from both the hyperplanes; thereby achieving the same effect. Such a distance via the data points will also make the frames loosely coupled. More formally, we propose to redefine $\text{dist}_{\mathbf{W}}^2$ as:

$$\text{dist}_{\mathbf{W}}^2(\mathbf{W}_1, \mathbf{W}_2, \mathbf{b}_1, \mathbf{b}_2 | \mathbf{x}) = \frac{1}{2} \sum_{j=1}^2 \left\| \mathbf{W}_j^T \mathbf{x} + \mathbf{b}_j \right\|_2^2, \quad (9)$$

where now we minimize the sum of the lengths of each \mathbf{x} after projecting on to the respective frames; thereby pulling both the frames closer to the data point. Using this definition of $\text{dist}_{\mathbf{W}}^2$, we formulate our *generalized one-class discriminative subspace* (GODS) classifier as:

$$\text{GODS} := \min_{\mathbf{W} \in \mathcal{S}_d^K, \mathbf{b}} F = \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^2 \left\| \mathbf{W}_j^T \mathbf{x}_i + \mathbf{b}_j \right\|_2^2 \quad (10)$$

$$+ \frac{\nu}{2n} \sum_i \left[\eta - \min(\mathbf{W}_1^T \mathbf{x}_i + \mathbf{b}_1) \right]_+^2 + \left[\eta + \max(\mathbf{W}_2^T \mathbf{x}_i + \mathbf{b}_2) \right]_+^2$$

See Figure 2(c) for a graphical illustration of this variant of our GODS objective.

4.3 Extensions to GODS Formulation

The technical development of GODS in the previous section assumes the input data is unit normalized, as otherwise the orthonormal frames for discriminating them may not be generally fruitful. Our other important assumption in the previous section –

that the hyperplanes in our discriminative decision parameters \mathbf{W} are orthonormal – can be restrictive as well. In this section, we provide extensions of our GODS framework that relax or remove these restrictions. In the following variants, we will assume $\lambda > 0$ to generically denote a penalty on the respective regularization.

4.3.1 Non-Compact Stiefel Manifold

A matrix $\mathbf{W} \in \mathbb{R}^{d \times K}$ with its K columns being linearly independent, however not unit normalized, belongs to the so-called non-compact Stiefel manifold (Absil et al. [26][Chapter 3]), denoted $\mathbb{R}_*^{d \times K}$, which is an open subspace of the Euclidean space $\mathbb{R}^{d \times K}$. One may represent such a manifold as a product manifold between a $d \times K$ Stiefel manifold and a $K \times 1$ Euclidean vector; i.e., $\mathbb{R}_*^{d \times K} = \mathcal{S}_d^K \times \mathbb{R}^K$. For example, if $\mathbf{W} \in \mathbb{R}_*^{d \times K}$, then $\mathbf{W} = \mathbf{Q} \text{diag}(\mathbf{r})$, where $\mathbf{Q} \in \mathcal{S}_d^K$, $\mathbf{r} \in \mathbb{R}^K$, and the i -th dimension $\mathbf{r}_i = \|\mathbf{W}_{:,i}\|_2$, the ℓ_2 norm of the i -th column in \mathbf{W} . For non-unit-norm input data, we can extend the GODS formulation in (10) using a non-compact Stiefel manifold as:

$$\text{GODS}_N := \min_{(\mathbf{Q}, \mathbf{r}) \in \mathcal{S}_d^K \times \mathbb{R}^K, b} F(\mathbf{Q} \text{diag}(\mathbf{r}), b) + \frac{\lambda}{2} \|\mathbf{r}\|_p, \quad (11)$$

where F is the objective in (10) and $\lambda > 0$ is a penalty on the ℓ_p -norm regularization over \mathbf{r} . Note that in (11), for brevity we assume $(\mathbf{Q}, \mathbf{r}) = \{(\mathbf{Q}_i, \mathbf{r}_i)\}_{i=1}^K$, i.e., it is technically a product of two non-compact Stiefel manifolds corresponding to the two discriminative frames. There is an additional advantage with the proposed subspace representation – it can allow automatic selection of the number of subspace components one may need. For example, with the ℓ_p regularization on \mathbf{r} , some of the dimensions in \mathbf{r} can go to zero (say using $p = 1$), thereby removing the respective subspace component from the final representation.

4.3.2 Oblique Manifold

We may also relax the orthogonality constraints on \mathbf{W} , however maintain their unit normality. A set of matrices $\mathcal{OB}_d^K = \{\mathbf{W} \in \mathbb{R}^{d \times K} : \text{diag}(\mathbf{W}^T \mathbf{W}) = \mathbf{I}_K\}$ forms a regular submanifold of the Euclidean manifold and is usually called an Oblique manifold [101] under the canonical inner product metric. This manifold is isometric to the product of K spheres $\times_1^K \mathcal{S}_d^1$. We can rewrite a variant of GODS with the optimization on \mathcal{OB}_d^K as:

$$\text{GODS}_O := \min_{\mathbf{W} \in \mathcal{OB}_d^K, b} F(\mathbf{W}, b) + \frac{\lambda}{2} \sum_{i=1}^2 \left\| \mathbf{W}_i^T \mathbf{W}_i - \mathbf{I}_K \right\|_F^2, \quad (12)$$

the last term *softly* controls the correlations among columns in \mathbf{W} .

4.3.3 Euclidean Manifold

Removing both the orthogonality and the unit norm constraints on the classifiers and the input data results in our most general form of the GODS formulation, that assumes \mathbf{W} belongs to the Euclidean manifold. We can write such a variant as:

$$\text{GODS}_E := \min_{\mathbf{W} \in \mathbb{R}^{d \times K}, b} F(\mathbf{W}, b) + \frac{\lambda}{2} \sum_{i=1}^2 \left\| \mathbf{W}_i^T \mathbf{W}_i - \mathbf{I}_K \right\|_F^2. \quad (13)$$

Similar to (12), the last term in GODS_E controls the correlations between the columns in \mathbf{W} ; a large λ will promote \mathbf{W} to be similar to the original GODS formulation using the Stiefel manifold in (10).

4.4 Kernelized One-class Discriminative Subspaces

While, the classifier as described in our GODS formulation can offer computationally efficient, yet non-linear decision functions in the input space, it may fail in situations when input data cannot be bounded using rectilinear coordinates. For example, when the outliers form dense regions inside the one-class, or when the normal data forms islands, are xor-shaped, or ring-shaped. This lends a kernelized variant of GODS scheme inevitable.

To derive kernelized GODS, we use our formulation⁴ in P_2 , however expand the min and max constraints in (6) and (7) via propagating the η to each of the K hyperplanes. Such a simplification allows for a direct application of the Lagrange multipliers to derive the dual. We also assume that there are no outliers in the one-class data provided. This allows us to simplify the expressions we derive.⁵ With these simplifications, we rewrite our modified P_2 as:

$$\min_{\mathbf{W} \in \mathcal{S}_d^K, \mathbf{b}} \frac{1}{2} \|\mathbf{W}_1 - \mathbf{W}_2\|_F^2 + \frac{1}{2} \|\mathbf{b}_1 - \mathbf{b}_2\|_2^2 \quad (14)$$

$$\text{s.t. } \eta - (\mathbf{W}_{1j}^T \mathbf{x}_i + \mathbf{b}_1) \leq 0, \quad \forall j \in [K], i \in [n]$$

$$\eta + (\mathbf{W}_{2j}^T \mathbf{x}_i + \mathbf{b}_2) \leq 0, \quad \forall j \in [K], i \in [n].$$

Using the fact that $\mathbf{W}^T \mathbf{W} = \mathbf{I}_K$, and using non-negative dual variables $\mathbf{Y}, \mathbf{Z} \in \mathbb{R}_+^{K \times n}$, we have the following Lagrangian formulation of (14):

$$\begin{aligned} L(\mathbf{W}, \mathbf{b}, \mathbf{Y}, \mathbf{Z}) := & -\text{Tr}(\mathbf{W}_1^T \mathbf{W}_2) + \frac{1}{2} \|\mathbf{b}_1 - \mathbf{b}_2\|_2^2 + \\ & \text{Tr}(\mathbf{Y}^T (\eta \mathbf{1}_{K \times n} - \mathbf{W}_1^T \mathbf{X} - \mathbf{b}_1 \mathbf{1}_n^T)) + \\ & \text{Tr}(\mathbf{Z}^T (\mathbf{W}_2^T \mathbf{X} + \mathbf{b}_2 \mathbf{1}_n^T + \eta \mathbf{1}_{K \times n})). \end{aligned} \quad (15)$$

A straightforward reduction provides the following dual:

$$\min_{\mathbf{Y}, \mathbf{Z} \in \mathbb{R}_+^{K \times n}} \frac{1}{2} \mathbf{1}_n^T \mathbf{Y}^T \mathbf{Y} \mathbf{1}_n + \text{Tr}(\mathbf{Y} \mathbb{K} \mathbf{Z}^T) - \eta \text{Tr}((\mathbf{Y} + \mathbf{Z})^T \mathbf{1}_{K \times n})$$

$$\text{s.t. } (\mathbf{Y} - \mathbf{Z}) \mathbf{1}_n = \mathbf{0} \quad (16)$$

$$\mathbf{Y} \mathbb{K} \mathbf{Y}^T = \mathbf{Z} \mathbb{K} \mathbf{Z}^T = \mathbf{I}_K, \quad (17)$$

where $\mathbb{K} = \mathbf{X}^T \mathbf{X}$ is a linear kernel, however could be replaced by any other positive definite kernel via the kernel trick. We call our formulation above as *kernelized one-class discriminative subspaces* (KODS). Recall that, for $\mathbb{K} \in \mathbb{R}^{n \times n} \succ 0$, the constraints in (17) pose the KODS objective on the generalized Stiefel manifold $\mathcal{G}_{\mathbb{K}}^{K \times n}$, formally defined as $\mathcal{G}_{\mathbb{K}}^{K \times n} = \{\Lambda \in \mathbb{R}^{K \times n} : \Lambda \mathbb{K} \Lambda^T = \mathbf{I}_K, \Lambda \succ 0\}$. However, there are two constraints in our objective that adds hurdle to directly using this manifold for optimization: (i) the null-space constraint in (16), and (ii) the requirement that the dual variables \mathbf{Y}, \mathbf{Z} are non-negative. Below, we present soft-constraints circumventing these challenges and derive an approximate KODS objective.

4.4.1 Approximate KODS Formulation

We avoid the null-space constraint in KODS via incorporating (16) as a soft-constraint into the KODS objective using a regularization penalty, $\lambda > 0$. To circumvent the non-negative constraints, we replace the dual variables \mathbf{Y}, \mathbf{Z} by their element-wise squares,

4. We attempted to use other GODS variants, however they resulted in objectives that seemed computationally expensive.

5. We also note that the use of slacks brings in additional constraints, making our optimization setup computationally difficult.

e.g., $\mathbf{Y} \in \mathbb{R}^{K \times n} : \mathbf{Y} \rightarrow (\mathbf{Y} \odot \mathbf{Y})$, while retaining $\mathbf{Y} \in \mathcal{G}_{\mathbb{K}}^{K \times n}$. Note that the latter heuristic has been used before, such as in approximating quadratic assignment problems [102]. With these changes, we provide our *approximate KODS formulation* as:

$$\begin{aligned} \min_{\mathbf{Y}, \mathbf{Z} \in \mathcal{G}_{\mathbb{K}}^{K \times n}} \mathcal{K}(\mathbf{Y}, \mathbf{Z}) = & \frac{1}{2} \mathbf{1}_n^T (\mathbf{Y} \odot \mathbf{Y})^T (\mathbf{Y} \odot \mathbf{Y}) \mathbf{1}_n \\ & + \text{Tr}((\mathbf{Y} \odot \mathbf{Y}) \mathbb{K} (\mathbf{Z} \odot \mathbf{Z})^T) \\ & - \eta \text{Tr}(((\mathbf{Y} \odot \mathbf{Y}) + (\mathbf{Z} \odot \mathbf{Z}))^T \mathbf{1}_{K \times n}) \\ & + \frac{\lambda}{2} \|((\mathbf{Y} \odot \mathbf{Y}) - (\mathbf{Z} \odot \mathbf{Z})) \mathbf{1}\|_2^2, \end{aligned} \quad (18)$$

where the last factor corresponds to (16). Note that we use the squared form only on the optimization variables, and not on the constraints, and thus our objective is still on the generalized Stiefel product manifold.

To derive the classification rules at test time (in the next section), we will need expressions for the primal variables in terms of the duals, which we provide below:

$$\mathbf{W}_1(\cdot) = (\mathbf{Z} \odot \mathbf{Z}) \mathbb{K}(\mathbf{X}, \cdot) \quad (19)$$

$$\mathbf{W}_2(\cdot) = -(\mathbf{Y} \odot \mathbf{Y}) \mathbb{K}(\mathbf{X}, \cdot) \quad (20)$$

$$\mathbf{b}_1 = \text{rowmax}(\eta - (\mathbf{Z} \odot \mathbf{Z}) \mathbb{K}) \quad (21)$$

$$\mathbf{b}_2 = \text{rowmin}(-\eta + (\mathbf{Y} \odot \mathbf{Y}) \mathbb{K}), \quad (22)$$

where rowmax and rowmin corresponds to the maximum and minimum values along the rows of the respective matrices.

5 ONE-CLASS CLASSIFICATION

During inference, we use the decision functions with the learned parameters to classify a given data point as in-class or out-of-class. Specifically, for a new data point \mathbf{x} , it is classified as in-class if the following criteria is met:

$$\min(\mathbf{W}_1(\mathbf{x}) + \mathbf{b}_1) \geq \eta \wedge \max(\mathbf{W}_2(\mathbf{x}) + \mathbf{b}_2) \leq -\eta, \quad (23)$$

where the variables \mathbf{W} and \mathbf{b} are either learned in the KODS formalism or the GODS. In case, we have access to a validation set consisting of in-class and out-of-class data (for which we know the class labels), then we may calibrate the threshold η to improve our decision rules. Specifically, suppose we have access to m such validation data points, denoted \mathcal{D}_v . Then, to estimate an updated threshold η' , we propose to compute the decision scores $v_l = \{\min(\mathbf{W}_1(\mathbf{x}) + \mathbf{b}_1)\}_{\mathbf{x} \in \mathcal{D}_v}$ and $v_u = \{\max(\mathbf{W}_2(\mathbf{x}) + \mathbf{b}_2)\}_{\mathbf{x} \in \mathcal{D}_v}$. Next, we apply K-Means (or spectral clustering) on v_l and v_u with $K = 2$ clusters. Suppose c_{lk} and c_{uk} ($k = 1, 2$) are the respective centroids for the two clustering problems; then we propose to update η' as the average of the smaller of the two centroids thresholded by η ; i.e.,

$$\Delta \eta = \frac{1}{2} ([\eta - \min(c_{l1}, c_{l2})]_+ - [\eta + \min(c_{u1}, c_{u2})]_+), \quad (24)$$

and use $\eta' = \eta + \Delta \eta$ to form the new decision rules in (23).

6 GODS OPTIMIZATION

In contrast to OC-SVM and SVDD, the GODS formulation in (10) is non-convex due to the orthogonality constraints on \mathbf{W}_1 and \mathbf{W}_2 .⁶ However, these constraints naturally impose a geometry

6. Note that the function $\max(0, \min(z))$ for z in some convex set is also non-convex.

on the solution space and in our case, puts optimization on the Stiefel manifold [103] – a Riemannian manifold characterizing the space of all orthogonal frames. There exist several schemes for geometric optimization over Riemannian manifolds (see [26] for a detailed survey) from which we use the Riemannian conjugate gradient (RCG) scheme in this paper, due to its stable and fast convergence. In the following, we review some essential components of the RCG scheme and provide the necessary formulae for using it to solve our objectives.

6.1 Riemannian Conjugate Gradient

Recall that the standard (Euclidean) conjugate gradient (CG) method [26][Sec.8.3] is a variant of the steepest descent method, however chooses its descent along directions conjugate to previous descent directions with respect to the parameters of the objective. Formally, suppose $F(\mathbf{W})$ represents our objective.⁷ Then, the CG method uses the following recurrence at the k -th iteration:

$$\mathbf{W}^k = \mathbf{W}^{k-1} + \lambda^{k-1} \alpha^{k-1}, \quad (25)$$

where λ is a suitable step-size (found using line-search) and $\alpha^{k-1} = -\text{grad } F(\mathbf{W}^{k-1}) + \mu^{k-1} \alpha^{k-2}$, where $\text{grad } F(\mathbf{W}^{k-1})$ defines the gradient of F at \mathbf{W}^{k-1} and α^{k-1} is a direction built over the current residual, which is conjugate to previous descent directions (see [26][pp.182]).

When \mathbf{W} belongs to a curved Riemannian manifold, we may use the same recurrence, however there are a few important differences from the Euclidean CG case, namely (i) we need to ensure that the updated point \mathbf{W}^k belongs to the manifold, (ii) there exists efficient vector transports⁸ for computing α^{k-1} , and (iii) the gradient grad is along tangent spaces to the manifold. For (i) and (ii), we may resort to computationally efficient retractions (using QR factorizations; see [26][Ex.4.1.2]) and vector transports [26][pp.182], respectively. For (iii), there exist standard ways that take as input a Euclidean gradient of the objective (i.e., assuming no manifold constraints exist), and maps them to the Riemannian gradients [26][Chap.3]. Specifically, for the Stiefel manifold, let $\nabla_{\mathbf{W}} F(\mathbf{W})$ define the Euclidean gradient of F (without the manifold constraints), then the Riemannian gradient is given by:

$$\text{grad } F(\mathbf{W}) = \nabla_{\mathbf{W}} F(\mathbf{W}) - \mathbf{W} \nabla_{\mathbf{W}} F(\mathbf{W})^{\top} \mathbf{W}. \quad (26)$$

The direction $\text{grad } F(\mathbf{W})$ corresponds to a curve along the manifold, descending along which ensures the optimization objective is decreased (atleast locally).

Now, getting back to our one-class objective, all we need to derive to use the RCG, is compute the Euclidean gradients $\nabla_{\mathbf{W}} F(\mathbf{W})$ of our objective in GODS with regard to the variables \mathbf{W}_j s. The other variables, such as the biases and slacks, belong to the Euclidean space and their gradients are straightforward. The expression for the Euclidean gradient of our objective with respect to the \mathbf{W} 's is given by:

7. The other optimization variable – \mathbf{b} , belongs to the Euclidean manifold, and thus $(\mathbf{W}, \mathbf{b}) \in \mathcal{S}_d^K \times \mathbb{R}^K$. However for brevity and focus, we omit these variables from our optimization discussion.

8. This is required for computing α_{k-1} that involves the sum of two terms in potentially different tangent spaces, which would need vector transport for moving between them; see [26][pp.182].

$$\frac{\partial F}{\partial \mathbf{W}_1} = \sum_{i=1}^n \mathbf{x}_i \left(\mathbf{W}_1^{\top} \mathbf{x}_i + \mathbf{b}_1 \right)^{\top} - \mathbf{Z}_{k_i^*} \left[\eta - \mathbf{W}_{1,k_i^*}^{\top} \mathbf{x}_i - b_1 \right]_{+}, \quad (27)$$

$$\frac{\partial F}{\partial \mathbf{W}_2} = \sum_{i=1}^n \mathbf{x}_i \left(\mathbf{W}_2^{\top} \mathbf{x}_i + \mathbf{b}_2 \right)^{\top} + \mathbf{Z}_{k_i^*} \left[\eta + \mathbf{W}_{2,k_i^*}^{\top} \mathbf{x}_i + b_2 \right]_{+}, \quad (28)$$

where $k_i^* \in [K]$ denotes the hyperplane index for the respective subspaces; $k_i^* = \arg \min_k (\mathbf{W}_1^{\top} \mathbf{x}_i + \mathbf{b}_1)$ for (27) and $k_i^* = \arg \max_k (\mathbf{W}_2^{\top} \mathbf{x}_i + \mathbf{b}_2)$ for (28). The variable $\mathbf{Z}_{k_i^*}$ is a $d \times K$ matrix with all zeros, except k_i^* -th column, which is set to \mathbf{x}_i .

6.2 KODS Optimization

In this section, we will derive the gradients for our approximate KODS formulation provided in (18). Similar to (26), the mapping from the Euclidean gradient to the Riemannian gradient for the generalized Stiefel manifold is provided in the following theorem.

Theorem 1. *For the optimization problem $\min_{\mathbf{U}} \mathcal{K}(\mathbf{U})$ s.t. $\mathbf{U} \mathbb{K} \mathbf{U}^{\top} = \mathbf{I}_K$, $\mathbb{K} \succ 0$, if $\nabla_{\mathbf{U}} \mathcal{K}(\mathbf{U})$ denotes the Euclidean gradient of $\mathcal{K}(\mathbf{U})$, then the Riemannian gradient under the canonical metric is given by:*

$$\text{grad } \mathcal{K}(\mathbf{U}) = \nabla_{\mathbf{U}} \mathcal{K}(\mathbf{U}) \mathbb{K}^{-1} - \mathbf{U} \nabla_{\mathbf{U}} \mathcal{K}(\mathbf{U})^{\top} \mathbf{U}. \quad (29)$$

Proof. The proof follows directly from the results in [24][Section 4.5]. \square

For the retraction of the iterates on to the manifold, we use the generalized polar decomposition [104] as suggested in [105]. As in the previous section, next we derive the expressions for the Euclidean gradients.

Proposition 1. *Let $f(\mathbf{Y})$ be a differentiable matrix function, then the matrix gradient $\nabla_{\mathbf{Y}^{\top}} f(\mathbf{Y}) = (\nabla_{\mathbf{Y}} f(\mathbf{Y}))^{\top}$.*

Lemma 1. *For matrices \mathbf{Y} , \mathbf{A} , and \mathbf{D} of appropriate sizes, if $f(\mathbf{Y}) = \text{Tr}((\mathbf{Y} \odot \mathbf{A}) \mathbf{D})$, then $\nabla_{\mathbf{Y}} f(\mathbf{Y}) = \mathbf{A} \odot \mathbf{D}^{\top}$.*

Proof. If $\mathbf{a}_{i:}$, $\mathbf{y}_{i:}$, $\mathbf{d}_{:j}$ represent the i -th row and j -th column of matrices \mathbf{A} , \mathbf{Y} , \mathbf{D} respectively, then

$$f(\mathbf{Y}) = \sum_i (\mathbf{a}_{i:} \odot \mathbf{y}_{i:})^{\top} \mathbf{d}_{:i} = \sum_i \mathbf{y}_{i:}^{\top} (\mathbf{a}_{i:} \odot \mathbf{d}_{:i}). \quad (30)$$

Then, the gradient w.r.t. \mathbf{y}_{ij} , i.e., $\nabla_{\mathbf{y}_{ij}} f(\mathbf{Y}) = \mathbf{a}_{ij} \odot \mathbf{d}_{ji}$, and we have the desired result. \square

Lemma 2. *Let $f(\mathbf{Y}) = \text{Tr}((\mathbf{Y} \odot \mathbf{Y})^{\top} (\mathbf{Y} \odot \mathbf{Y}) \mathbf{D})$, where \mathbf{D} is a symmetric matrix. Then, $\nabla_{\mathbf{Y}} f(\mathbf{Y}) = 4\mathbf{Y} \odot (\mathbf{Y} \odot \mathbf{Y}) \mathbf{D}$.*

Proof. To simplify the notation, let us use $\mathbf{A} = (\mathbf{Y} \odot \mathbf{Y})$, then using Proposition 1, and applying chain-rule:

$$\begin{aligned} \nabla_{\mathbf{Y}} f(\mathbf{Y}) &= 2 \left(\nabla_{\bullet} \text{Tr} \left((\bullet^{\top} \odot \mathbf{Y}^{\top}) \mathbf{A} \mathbf{D} \right) \right)^{\top} \\ &\quad + 2 \nabla_{\bullet} \text{Tr} \left(\mathbf{A}^{\top} (\bullet \odot \mathbf{Y}) \mathbf{D} \right) \\ &= 2 \left(\mathbf{Y}^{\top} \odot (\mathbf{A} \mathbf{D})^{\top} \right)^{\top} + 2 \mathbf{Y} \odot (\mathbf{D} \mathbf{A}^{\top})^{\top} \quad (31) \\ &= 2 \mathbf{Y} \odot (\mathbf{A} \mathbf{D}) + 2 \mathbf{Y} \odot (\mathbf{A} \mathbf{D}^{\top}), \end{aligned}$$

where we used Lemma 1 to obtain (31). Using the symmetry of \mathbf{D} , we have the result. \square



Fig. 3: Frames from our Dash-Cam-Pose dataset. The left-top frame has poses in-position (one-class), while the rest of the frames are from videos labeled out-of-position.

Theorem 2. Let $\mathbf{E}_n = \mathbf{1}_n \mathbf{1}_n^\top$, $\mathbf{Y}^2 = \mathbf{Y} \odot \mathbf{Y}$, and $\mathbf{Z}^2 = \mathbf{Z} \odot \mathbf{Z}$, the Euclidean gradient of $\mathcal{K}(\mathbf{Y}, \mathbf{Z})$ in (18) is:

$$\begin{aligned} \nabla_{\mathbf{Z}} \mathcal{K}(\mathbf{Y}, \mathbf{Z}) &= 2\lambda \mathbf{Z} \odot \mathbf{Z}^2 \mathbf{E}_n + \mathbf{Z} \odot \mathbf{Y}^2 [2\mathbb{K} - \lambda \mathbf{E}_n] - 2\eta \mathbf{Z} \\ \nabla_{\mathbf{Y}} \mathcal{K}(\mathbf{Y}, \mathbf{Z}) &= \gamma \mathbf{Y} \odot \mathbf{Y}^2 \mathbf{E}_n + \mathbf{Y} \odot \mathbf{Z}^2 [2\mathbb{K} - \lambda \mathbf{E}_n] - 2\eta \mathbf{Y}, \end{aligned}$$

where $\gamma = 2 + 2\lambda$.

Proof. The result directly follows by applying Lemma 1 and Lemma 2 to the formulation in (18). \square

6.3 Optimization Initialization

Due to the non-convexity of our objective, there could be multiple local solutions. To this end, we resort to the following initialization of our optimization variables, which we found to be empirically beneficial. Specifically, for the GODS optimization, we first sort all the training points based on their Euclidean distances from the origin. Next, we randomly select a suitable number ($3 \times \#$ hyperplanes in our experiments) of such sorted points near and far from the origin, compute a compact singular value decomposition (thin-SVD) of these points, and initialize the GODS subspaces using these orthonormal frames from the SVD. The intercepts (\mathbf{b}) are initialized to zero. For KODS, we initialize the dual variables as $\frac{1}{nK}$, where K is the number of hyperplanes.

7 EXPERIMENTS

In this section, we provide experiments demonstrating the performance of our proposed schemes on several one-class tasks. We will introduce these tasks and the associated datasets briefly next along with detailing the data features used.

7.1 Dash-Cam-Pose Dataset

Out-of-position (OOP) human pose detection is an important problem with regard to the safety of passengers in a vehicle. While, there are public datasets for human pose estimation, they are usually annotated for generic pose estimation tasks, and neither do they contain any in-vehicle poses as captured by a dashboard camera, nor are they annotated for pose anomalies. To this end, we collected 104 videos, each 20-30 min long, from the Internet (including Youtube, Shutterstock, and Hollywood road movies). As these videos were originally recorded for diverse reasons, there




Fig. 4: Some examples from JHMDB (first column), UCF-Crime (second column), and USCD Ped2 (third column) datasets, with respective categories.

are significant shifts in camera angles, perspectives, locations of the camera, scene changes, etc. We encourage the interested reader to refer to [32] for more details on this dataset and its collection. Next, we manually selected clips from these videos that are found interesting for our task. To extract as many clips as possible from these videos, we first segmented each video into three second clips at 30fps, which resulted in approximately 7000 clips. Next, we selected only those clips where the camera is approximately placed on the dashboard looking inwards, which amounted to 4,875 clips, totalling 4.06 hours. We annotated each clip with a weak binary label based on the poses of humans in the front seat. Specifically, if all the front-seat humans (passengers and the driver) are seated in-position, the clip was given a positive label, while if any human is seated OOP (based on [106], [107]) for the entire 3s, the clip was labeled as negative. Next, we used Open Pose [108] on each clip to extract a sequence of poses for every person. Our final Dash-Cam-Pose dataset consists of 4875 short videos, 1.06 million poses, of which 310,996 are OOP.

We explore two pose-sequence representations for this task: (i) a simple bag-of-words (BoW) model, and (ii) using a Temporal Convolutional Network (TCN) [109] consisting of residual units with 1D convolutional layers capturing both local and global information via convolutions for each joint across time. For BoW, we use 1024 pose centroids computed using K-Means clustering. For the TCN, we use the following procedure. The poses from each person in each frame are vectorized and stacked into the temporal dimension. For each pose thus passed through TCN, we extract features from the last pooling layer, using a model pre-trained on the NTU-RGBD dataset [110] (for 3D skeleton action recognition) to produce 256-D features for every clip. As our pre-trained TCN model takes 3D poses as input, we pad our Dash-Cam-Poses with zeros in the third dimension.

We use a four-fold cross-validation for evaluation on Dash-Cam-Pose. Specifically, we divide the entire dataset into four non-overlapping splits, each split consisting of approximately 1/4-th the dataset, of which roughly 2/3rd's are labeled as positive (in-pose) and the rest as OOP. We use only the positive data in each split to train our one-class models. Once the models are trained, we evaluate on the held out split. For every embedded-pose feature, we use the binary classification accuracy against the ground truth. The evaluation is repeated on all the four splits and the performance averaged.

7.2 Public Datasets

JHMDB dataset: is a video action recognition dataset [31] consisting of 968 clips with 21 classes (see Figure 4 for example frames). To adapt the dataset for a one-class evaluation, we use a one-versus-rest strategy by choosing sequences from an action class as “normal” while those from the rest 20 classes are treated as “abnormal”. To evaluate the performance on the entire dataset, we cycle over the 21 classes, and the scores are averaged. For representing the frames, we use an ImageNet pre-trained VGG-16 model and extract features from the ‘fc-6’ layer (4096-D).

UCF-Crime dataset: is the largest publicly available real-world anomaly detection dataset [29], consisting of 1900 surveillance videos and 13 crime categories such as *fighting*, *robbery*, etc. and several “normal” activities, such as the daily walking, running and driving. Illustrative video frames from this dataset and their class labels are shown in Figure 4. To encode the videos, we use the state-of-the-art Inflated-3D (I3D) neural network [111]. Specifically, video frames from non-overlapping sliding windows (8 frames each) are passed through the I3D network; features are extracted from the ‘Mix_5c’ network layer, that are then reshaped to 2048-D vectors. For anomaly detections on the test set, we first map back the features classified as anomalies by our scheme to the frame-level and apply the official evaluation metrics [29].

UCSD Ped2 dataset: contains 16 videos in the training and 12 videos in the test set. There are 12 abnormal events in the test videos, such as the *Biker*, *Cart*, *Skater*, etc. To encode the video data, we apply the deep autoencoder with causal 3D convolutions [112] trained to minimize the reconstruction loss. We extract features from the bottleneck layer of this model to be input to our algorithm. As the videos can be of arbitrary length, the pipeline is trained on clips from temporal sliding windows with a stride of one and consisting of 16 frames. As anomalous events are labeled frame-wise in this dataset, we use the averaged clip-level predictions within a window as the prediction for the center frame in that window. We use the evaluation metrics on these frame-level predictions similar to [9], [17], [19].

7.3 Experimental Setup

Before using the above features in our algorithms, we found that it is beneficial to unit-normalize them. However, we do report results without such normalization on other datasets in Section 7.7. These scaled features are then used in our GODS formulations, the optimization schemes for which are implemented using ManOpt [105] and PyManOpt [113]. We use the conjugate gradient scheme for optimization, which typically converges in about 200 iterations. We initialize the iterates using the approach described in Section 6.3. The hyper-parameters in our models are chosen via cross-validation, and the sensitivities of these parameters are evaluated in the next section. We use regularization constants $\nu = 1$ and $\lambda = 1$. We use the inference criteria described in Section 5 for classifying a test point as in-class or an anomaly.

7.4 Evaluation Metrics

On the UCF-Crime dataset, we follow the official evaluation protocol, reporting AUC as well as the false alarm rate. For other datasets, we use the $F1$ score to reflect the sensitivity and accuracy of our classification models. As the datasets we use - especially the Dash-Cam-Pose - are imbalanced across the two classes, having a single performance metric over the entire dataset may fail to characterize the quality of the discrimination for each

class separately, which is of primary importance for the one-class task. To this end, we also report True Negative Rate $TNR = \frac{TN}{N}$, Negative Predictive Value $NPV = \frac{TN}{TN+FN}$, and $\overline{F1} = \frac{2 \times TNR \times NPV}{TNR + NPV}$, alongside standard $F1$ scores. We will use $\overline{F1}$ on the Dash-Cam-Pose dataset and $F1$ score on other datasets. Informally, $\overline{F1}$ is the same as $F1$ with the positive and negative categories switched.

7.5 Ablative Studies

Synthetic Experiments: To gain insights into the inner workings of our schemes, we present results on several 2D synthetic toy datasets. In Figure 1(a)–1(c), we show three plots with 100 points distributed as (i) Gaussian and (ii) some arbitrary distribution⁹. We show the BODS hyperplanes in the Figure 1(a), and the GODS 2D subspaces in Figures 1(b), 1(c) with the hyperplanes belonging to each subspace shown in same color. As the plots show, our models are able to orient the subspaces such that they confine the data within a minimal volume. In Figure 1(d), we show 300 data points (black dots) distributed along a 2D ring, a situation when a rectilinear GODS may fail (as the inner circle does not contain the one-class). As seen in Figures 1(e) and 1(f), the two kernelized KODS hyperplanes capture the outer and inner decision regions separately, and their combined decision region is able to capture the ring structure of the input data, as seen in Figure 1(d). In Figure 1(g), we plot the decision surfaces for 3D data points.

Choice of Manifold and Initialization. As described in Section 4.3, our GODS algorithm may assume several optimization manifolds based on the type of regularization used between the classifier hyperplanes. In the Figure 5, we evaluate three different manifold choices, namely (i) Stiefel manifold, (ii) oblique manifold, and (iii) Euclidean manifold. We also evaluate three different hyperplane initialization strategies: (i) random, (ii) SVD (as described in Section 6.3), and (iii) mean of the data features. From the Figure, it can be seen that the Stiefel manifold and SVD initialization works best compared to oblique and Euclidean manifolds and against random or mean initializations, consistently on the three datasets.

Choice of Normalization and soft-orthogonality. In Table 1, we investigate the effectiveness of the L2 normalization and soft-orthogonality proposed in Equation (12) and (13). For fair comparison, we implement the GODS algorithm with Stiefel, non-compact Stiefel, Euclidean and Oblique manifolds and evaluate on three different datasets (in Dash-Cam-Pose dataset, we use BOW and TCN features as input). From the experimental result, it is clear that the L2 norm helps to maintain a good performance in the GODS algorithm (verifying our assumptions in Section 4.1.1). Thus, we will apply L2 norm by default in our following experiments.

Sensitivity of Margin η . The hyperparameter η decides the support margin between the hyperplanes and the one-class data. In Figure 6, we analyze the performance sensitivity against changes in η on the JHMDB, UCF-Crime, UCSD-Ped2, and the Dash-Cam-Pose datasets. To ensure the learning will not ignore the changes in η , we increased the regularization constants on the two terms involving η in (4). On both datasets, the TPR (true positive rate) increases for increasing η , while the TNR decreases with a higher value of η . This is because the distances between the two orthonormal frames ($\mathbf{W}_1, \mathbf{W}_2$) may become larger to satisfy the

9. The data follows the formula $f(x) = \sqrt{x} * (x + \text{sign}(\text{randn}) * \text{rand})$, where randn and rand are standard MATLAB functions.

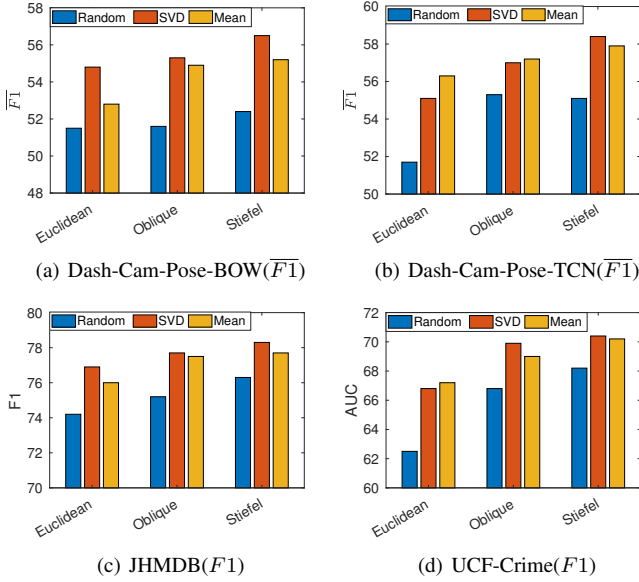


Fig. 5: Performance of GODS with $F1$, $\overline{F1}$ and AUC for different initialization methods and manifold assumptions.

TABLE 1: Comparisons of GODS variants on vision datasets. We compare: (i) GODS (Eq. (10)) using the Stiefel manifold, (ii) $GODS_N$ using the non-compact Stiefel (Eq. (11)), (iii) the Euclidean (Eq. (13)), and (iv) the oblique manifolds (Eq. (12)). We compare under (i) ℓ_2 unit-normalization of inputs and (ii) C , under soft-orthogonality (Eqs. (12), (13)). We report $F1$, $\overline{F1}$ and AUC scores (in %) for JHMDB, Dash-Cam-Pose (DCP), and UCF-Crime datasets.

	Type	DCP-BOW	DCP-TCN	JHMDB	UCF-Crime
GODS	ℓ_2	56.3	58.4	77.7	70.0
	$\neq \ell_2$	53.8	56.6	75.2	68.9
	$\neq \ell_2 + C$	53.5	55.4	75.1	67.3
	$\ell_2 + C$	54.9	58.5	77.5	70.3
$GODS_N$	ℓ_2	53.4	57.6	75.2	68.6
	$\neq \ell_2$	52.8	55.9	74.1	68.1
	$\neq \ell_2 + C$	52.4	54.3	72.9	67.7
	$\ell_2 + C$	53.3	56.5	74.8	69.1
$GODS_E$	ℓ_2	55.1	55.8	76.7	66.8
	$\neq \ell_2$	53.9	54.1	74.7	64.6
	$\neq \ell_2 + C$	53.2	53.7	73.6	63.4
	$\ell_2 + C$	54.8	55.1	74.9	64.8
$GODS_O$	ℓ_2	55.5	56.8	77.4	66.7
	$\neq \ell_2$	54.3	55.2	76.1	64.9
	$\neq \ell_2 + C$	53.8	54.5	75.4	64.8
	$\ell_2 + C$	54.8	55.9	76.8	66.2

new margin constraints imposed by η . Thus, more points will be included between the two frames and classified as positive. As $F1$ relies on the classification sensitivity of the positive data, while the $\overline{F1}$ relies on the negative classifications, they show opposite trends. Observing these trends, we fix $\eta = 0.3$ in our experiments. **Number of Hyperplanes K .** In Figure 7, we plot the influence of increasing number of hyperplanes on our four datasets. We find that after a certain number of hyperplanes, the performance saturates, which is expected, and suggests that more hyperplanes might lead to overfitting to the positive class. We also find that the TCN embedding is significantly better than the BoW model (by nearly 3%) on the Dash-Cam-Pose dataset when using our proposed methods. Surprisingly, S-SVDD is found to perform quite inferior against ours; note that this scheme learns a low-dimensional subspace to project the data to (as in PCA), and

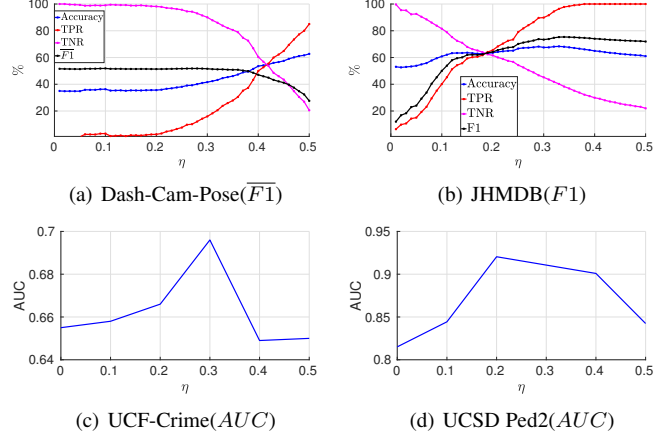


Fig. 6: Performance of BODS with $F1$ and $\overline{F1}$ for increasing η .

applies SVDD on this subspace. We believe, these subspaces perhaps are common to the negative points as well that it cannot be suitably discriminated, leading to poor performance. We make a similar observation on the other datasets as well.

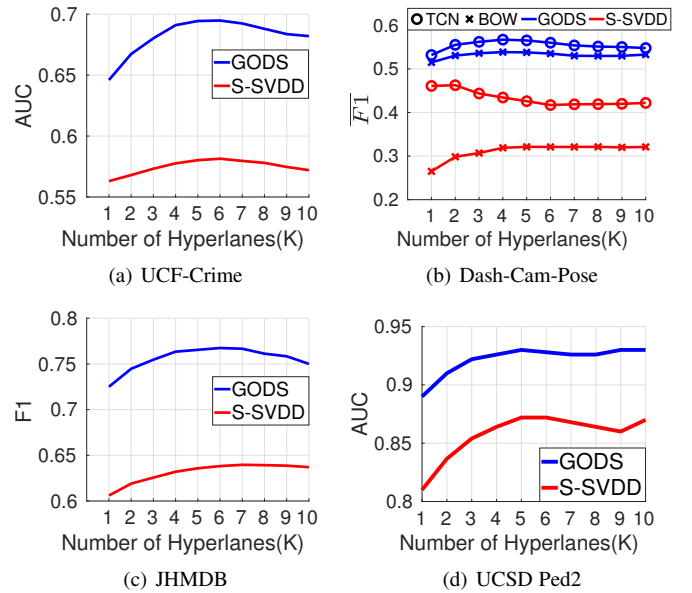


Fig. 7: Performance of GODS for an increasing number of subspaces.

Kernel Choices and Number of Subspaces in K . In Figure 8, we demonstrate the performance of KODS on the datasets for various choices of the embedding kernels, and also when increasing the number of Hilbert space classifiers K for each choice of the kernel feature map on every dataset we use. Specifically, we experiment with linear, RBF, and polynomial kernels on JHMDB, UCF-Crime, UCSD Ped2, and Dash-Cam-Pose (with TCN) datasets, while use the Chi-square [114] and Histogram Intersection kernels [115] on the Dash-Cam-Pose dataset with the Bag-of-Words features. We set $\sigma = 0.1$ for the bandwidth in the RBF kernel, polynomial kernel degree is set to 3, and use 1024 words in the Bag-of-Words representation. In Figure 8, we see that the performance saturates with increasing number of hyperplanes. The RBF kernel seems to work better on the JHMDB and UCF-Crime datasets, while the polynomial kernel demonstrates higher

performances on the UCSD-Ped2 and the Dash-Cam-Pose datasets (with TCN). For the Bag-of-Words features on the Dash-Cam-Pose dataset, the Chi-square kernel shows better performance than the Histogram Intersection kernel.

Empirical Convergence. In Figures 9(a) and 9(b), we show the empirical convergences of our GODS algorithm on the JHMDB dataset using the original GODS formulation in (10). We show the convergence in the objective value as well as the magnitude of the Riemannian gradients. As is clear, our algorithm convergences in about 200 iterations on this dataset. We repeat this experiment on the KODS formulation (18) using different kernel maps. As is seen from Figures 9(c) and 9(d), while the convergence is slower compared to that in GODS – perhaps due to our approximations – it does converge suitably for appropriate kernel choices.

Running Time. In the Figure 9(e), we demonstrate the time taken for training our different models. For this analysis, we use an Intel i7-6800K 3.4GHz CPU with 6 cores. We implement the different algorithms in the Matlab, run it on the same data, and record the training time with an increasing number of training samples. For GODS, KODS, and S-SVDD, we use 3 hyperplanes in the subspaces. It can be seen that the GODS, BODS, and KODS are not substantially more computationally expensive against prior methods, while remaining empirically superior (Table 2).

7.6 State-of-the-Art Comparisons

In Table 2, we compare our variants to the state-of-the-art methods. As alluded to earlier, for our Dash-Cam-Pose dataset, as its positive and negative classes are imbalanced, we resort to reporting the $\overline{F1}$ score on the test set. From the table, our variants are seen to outperform prior methods by a significant margin; especially our GODS and KODS schemes demonstrate the best performances on different tasks. For example, using TCN, KODS outperforms other kernelized prior variants by over 20%. Similarly, on the JHMDB dataset, both GODS and KODS are better than the next best kernel-based method (K-OC-SVM) by about 20%, and improves the classification accuracy by over 30%. Overall, the experiments clearly substantiate the performance benefits afforded by our methods on the one-class task.

In Table 3-left, we present results against the state of the art on the UCF-Crime dataset using the AUC metric and false alarm rates; we use the standard threshold of 50%. While, our results are lower than [29] by 4% in AUC and 0.2 larger in false alarm rate, their problem setup is completely different from ours in that they use weakly labeled abnormal videos as well in their training, which we do not use and which as per definition is not a one-class problem. Thus, our results are incomparable to theirs. Against other methods on this dataset, our schemes are about 5-10% better. In the Table 3-right, we provide the performance (AUC) on the UCSD Ped2 dataset. Compared with the recent state-of-the-art methods, both GODS and KODS achieves similar performances using the 3D autoencoder features. Among these methods, Luo et al. [68] and Liu et al. [19] propose ConvLSTM-AE and S-RNN respectively, which rely on the recurrent neural networks, that might be hard to train. Abati et al. [17] also uses the 3D autoencoder features similar to ours, but also employs additional constraints for building the Conditional Probability Density (CPD), which is more expensive compared to our solution. We also note that Table 3-right lists prior methods that use deep learning models, such as the ConvLSTM autoencoder [68], Stacked-RNN [19], and GANs [116]; our GODS variants offer competitive performances against them.

TABLE 2: Average performances on the Dash-Cam-Pose and JHMDB datasets. Dash-Cam-Pose uses the $\overline{F1}$ score while JHMDB uses $F1$ score as evaluation metric (classification accuracy is shown in the brackets). K-OC-SVM and K-SVDD are the RBF kernelized variants.

Method	CarPose_BOW	CarPose_TCN	JHMDB
OC-SVM [13]	0.167 (0.517)	0.279(0.527)	0.301 (0.568)
SVDD [16]	0.448 (0.489)	0.477(0.482)	0.407 (0.566)
K-OC-SVM [13]	0.327 (0.495)	0.361(0.491)	0.562 (0.412)
K-SVDD [16]	0.476 (0.477)	0.489 (0.505)	0.209 (0.441)
K-PCA [84]	0.145 (0.502)	0.258 (0.492)	0.245 (0.557)
Slab-SVM [89]	0.468 (0.568)	0.498 (0.577)	0.643 (0.637)
KNFST [87]	0.345 (0.487)	0.368 (0.496)	0.667 (0.501)
KNN [117]	0.232 (0.475)	0.276 (0.488)	0.643 (0.492)
LS-OSVM [14]	0.234 (0.440)	0.246(0.460)	0.663(0.582)
S-SVDD [81]	0.325 (0.490)	0.464 (0.500)	0.642 (0.498)
BODS	0.523 (0.582)	0.532 (0.579)	0.725 (0.714)
GODS	0.563 (0.629)	0.584 (0.601)	0.777 (0.752)
KODS	0.596 (0.642)	0.664 (0.604)	0.785 (0.726)

TABLE 3: Performances on UCF-Crime dataset (left) and UCSD Ped2 dataset (right). *Setup is different.

UCF Crime Dataset			UCSD Ped2 Dataset	
Method	AUC	FAR	Method	AUC
Random	0.50	-	ConvLSTM-AE. [68]	0.88
Hasan et al. [20]	0.51	27.2	GANs [116]	0.88
Lu et al. [56]	0.66	3.1	S-RNN [19]	0.92
*Waqas et al. [29]	0.75	1.9	Autoregression [17]	0.95
Sohrab et al. [81]	0.59	10.5	FFP+MC, Liu et al. [9]	0.95
BODS	0.68	2.7	BODS	0.91
GODS	0.70	2.1	GODS	0.93
KODS	0.71	2.1	KODS	0.95

7.7 Performance on UCI datasets

As the reader might acknowledge, the algorithms proposed in this paper are not specialized to only computer vision datasets, but could be applied for the anomaly detection task on any data mining, machine learning, or robotics task. To this end, in Table 4, we evaluate GODS and KODS on five datasets downloaded from UCI datasets¹⁰ and TU delft pattern recognition lab website¹¹. These datasets are: (i) sonar, the task in which is to discriminate between sonar signals bounced off a metal cylinder and those bounced off a roughly cylindrical rock, (ii) the Delft pump dataset, the task in which is to detect abnormal condition of a submersible pump, (iii) the Scale dataset is to classify the balance scale tip to the right, tip to the left, or be balanced, (iv) the Haberman’s survival dataset records the survival of patients who had undergone surgery for breast cancer, and (v) the Banknote dataset is to detect if the feature from an image passing the evaluation of an authentication procedure for banknotes.

We follow the evaluation protocol in the recent paper [81] for all the datasets and compare the performances to those reported in that paper. Specifically, the evaluation uses a split of 70/30 for the positive class; the model training is performed on the 70%, and tested on the remaining 30% positive class data and the negative (anomalous) data (which is not used in training). We repeat the split in the positive class five times and report the average performance on the five trials. For datasets having more than two classes, we pick one class as positive, while the remaining as negative, and follow the same protocol as above. In Table 4, we report the performances of GODS and KODS against those reported in [81]. In the table, N represents the number of samples in the dataset, D denotes the feature dimension of

10. <http://archive.ics.uci.edu/ml/index.php>

11. <http://homepage.tudelft.nl/n9d04/occ/index.html>




Fig. 8: Performance of KODS in different kernel type on various datasets for an increasing number of subspaces.




Fig. 9: (a–d) show optimization convergence on JHMDB dataset. (a-b) GODS (10) against the choice of optimization schemes, (c-d) KODS (18) under different kernels (using CG optimizer). (e) compares running time.

TABLE 4: Performances on UCI datasets. N is number of samples, D is the feature dimension and T is the target (positive) class.

No.	Dataset	N	D	T	OC-SVM [81]	K-OC-SVM	SVDD [81]	S-SVDD [81]	K-SVDD	GODS	KODS
1	Sonar	208	60	Mines	53.5%	56.5%	62.5%	63.8%	60.9%	71.6%	72.5%
2	Pump	1500	64	Normal	63.2%	60.1%	84.6%	85.7%	83.6%	87.6%	88.2%
3	Scale	625	4	Left	68.8%	67.6%	70.3%	90.7%	73.4%	92.3%	92.5%
4	Survival	306	3	Survived	64.4%	74.3%	83.4%	84.1%	83.5%	87.6%	88.1%
5	Banknote	1372	5	No	65.7%	70.0%	76.4%	90.8%	80.4%	94.7%	95.8%

TABLE 5: Comparisons of GODS variants on UCI datasets. We compare: (i) GODS (10) using the Stiefel manifold, (ii) $GODS_N$ using the non-compact Stiefel (11), (iii) the Euclidean (13), and (iv) the oblique manifolds (12). We compare under (i) ℓ_2 unit-normalization of inputs and (ii) C , under soft-orthogonality ((12), (13)). We report F1 scores (in %) and standard deviations over 5 trials.

	Type	Sonar	Pump	Scale	Survival	Banknote
GODS	ℓ_2	71.6±2.6	87.6±0.6	89.2±5.4	87.6±1.9	94.7±3.8
	$\neq \ell_2$	69.2±3.9	85.5±0.9	92.3±4.5	86.5±8.9	90.5±2.1
	$\neq \ell_2 + C$	68.3±4.1	85.2±0.8	92.2±4.1	87.0±9.5	90.4±2.1
	$\ell_2 + C$	71.6±3.7	87.0±0.6	88.7±3.6	87.8±1.8	95.1±1.7
$GODS_N$	ℓ_2	69.2±8.6	86.7±7.1	84.8±7.0	85.3±0.3	90.9±6.2
	$\neq \ell_2$	68.7±1.5	85.9±5.1	89.7±9.1	82.9±4.8	82.4±2.6
	$\neq \ell_2 + C$	66.9±2.2	85.7±4.1	89.5±9.5	85.2±7.2	88.4±4.5
	$\ell_2 + C$	69.1±5.1	86.4±0.6	86.9±4.2	86.3±7.8	90.9±6.4
$GODS_E$	ℓ_2	66.7±2.3	85.9±0.3	79.8±5.1	84.8±1.7	93.8±6.7
	$\neq \ell_2$	68.1±4.5	85.8±0.6	94.4±6.5	85.9±1.1	90.6±3.0
	$\neq \ell_2 + C$	68.2±3.0	85.9±0.6	95.0±2.7	86.5±7.5	91.4±2.2
	$\ell_2 + C$	68.7±2.4	86.2±0.2	81.0±3.9	86.5±1.2	94.1±1.5
$GODS_O$	ℓ_2	69.1±3.5	87.0±0.6	81.6±3.7	85.7±6.1	94.0±1.2
	$\neq \ell_2$	69.3±3.5	84.6±1.2	97.1±3.6	85.5±1.5	94.1±3.9
	$\neq \ell_2 + C$	70.5±4.6	86.5±0.7	97.1±2.9	85.9±8.7	95.3±3.2
	$\ell_2 + C$	70.6±4.0	87.3±0.7	82.3±4.3	85.9±3.3	95.6±3.1

each sample, and T is the target class picked as positive (same as in [81]). We use three hyperplanes in the GODS subspaces, and set the sensitivity margin $\eta = 0.3$ in both GODS and KODS. For KODS, we use polynomial kernel with degree as 3. As is clear from the table, we outperform the previous state-of-the-art results on all datasets. Specifically, GODS is substantially better than the previous best method S-SVDD by 2–8%, and the KODS is even better by 1–2%. This is because our GODS algorithm better characterizes the data distribution from positive classes and

thus producing higher cost for anomalies during the inference. For KODS, the kernel embedding would further bring advantages in learning the decision regions better fitting the normal samples.

In the Table 5, we evaluate the various extensions of GODS as described in Section 4.3. Through these experiments, we evaluate the impact of unit-normalization on the data inputs and the orthogonality assumptions on the hyperplanes, and analyze the adequacy of each variant when such assumptions may not be relevant. In the Table 5, each column contains the results for one dataset while the four rows in one column are the result for one variant of GODS. From top to the bottom, we have 4 settings relaxing different constraints; they are: 1) unit normalization applied on the inputs, denoted ℓ_2 , 2) no unit norm is enforced, $\neq \ell_2$, 3) $\neq \ell_2$ but with soft-orthogonality constraints (C) as described in (12), (13), and 4) ℓ_2 norm and C are used together. We also report the standard deviations associated with each experiment over the five trials.

From the experimental results, it is found that the orthogonal constraint is generally helpful, whenever applicable. For instance, the results in the third and fourth row in Euclidean and Oblique manifolds are better than the ones without orthogonal constraints by up to 2%. In terms of the ℓ_2 norm constraints, it depends on the nature of the data points. For example, in the Scale dataset, each dimension of the data captures the presence of some semantic attribute and thus ℓ_2 norm may not make much sense on them. However, for the vision datasets or the other four UCI datasets, the feature normalization could bound the data allowing the one-class model to better capture the distribution.

8 CONCLUSIONS

In this paper, we presented a novel one-class learning formulation – called GODS – using pairs of complementary classifiers; these classifiers are oriented so as to circumscribe the data density within a rectilinear space of minimal volume. We explored variants of GODS via relaxing the various constraints in our problem setup, as well as introducing kernel feature maps. Due to the orthonormality we impose on the classifiers, our objectives are non-convex, and solving for which we resorted to Riemannian optimization frameworks on the Stiefel manifold and its variants. We presented experiments on a diverse set of anomaly detection tasks, demonstrating state-of-the-art performances. We further analyzed the generalizability of our framework to non-vision data by presenting experiments on five UCI datasets; our results outperforming prior baselines by significant margins.

An potential direction to extend this work is perhaps to use more than two classifiers in the GODS framework. While, we experimented with a variant of this idea using multiple orthonormal frames, its performance was poor. We presume this inferior performance is perhaps due to the lack of appropriate regularizations across the classifiers and the absence of suitable complementarity conditions between the classifiers and the data. We plan to pursue this research direction in a future paper.

Further, there are several aspects of our scheme that needs rigorous treatment. For example, deriving generalization bounds on GODS is one such. Analysis of the representation complexity within a computational learning theory framework is yet another direction. An analysis of our optimization landscape is a direction that could help better initialize our schemes. Extending our framework as a module within an end-to-end neural network is an interesting direction as well.

Acknowledgements. We thank the anonymous reviewers whose insightful questions have helped us significantly improve the quality of this work.

REFERENCES

- [1] V. Chandala, A. Banerjee, and V. Kumar, “Anomaly detection: A survey, ACM computing surveys,” *University of Minnesota*, 2009.
- [2] C. M. Bishop, “Novelty detection and neural network validation,” *VISSP*, vol. 141, no. 4, pp. 217–222, 1994.
- [3] G. Ritter and M. T. Gallegos, “Outliers in statistical pattern recognition and an application to automatic chromosome classification,” *Pattern Recognition Letters*, vol. 18, no. 6, pp. 525–539, 1997.
- [4] A. B. Gardner, A. M. Krieger, G. Vachtsevanos, and B. Litt, “One-class novelty detection for seizure analysis from intracranial eeg,” *JMLR*, vol. 7, no. Jun, pp. 1025–1044, 2006.
- [5] S. S. Khan and M. G. Madden, “A survey of recent trends in one class classification,” in *Irish conference on artificial intelligence and cognitive science*. Springer, 2009, pp. 188–197.
- [6] M. Sabokrou, M. Khalooei, M. Fathy, and E. Adeli, “Adversarially learned one-class classifier for novelty detection,” in *CVPR*, 2018, pp. 3379–3388.
- [7] A. Del Giorno, J. A. Bagnell, and M. Hebert, “A discriminative framework for anomaly detection in large videos,” in *ECCV*, 2016.
- [8] V. Saligrama and Z. Chen, “Video anomaly detection based on local statistical aggregates,” in *CVPR*. IEEE, 2012, pp. 2112–2119.
- [9] W. Liu, W. Luo, D. Lian, and S. Gao, “Future frame prediction for anomaly detection—a new baseline,” in *CVPR*, 2018, pp. 6536–6545.
- [10] A. Buades, B. Coll, and J.-M. Morel, “A non-local algorithm for image denoising,” in *CVPR*. IEEE, 2005, pp. 60–65.
- [11] Y. Xia, X. Cao, F. Wen, G. Hua, and J. Sun, “Learning discriminative reconstructions for unsupervised outlier removal,” in *ICCV*, 2015, pp. 1511–1519.
- [12] C. You, D. P. Robinson, and R. Vidal, “Provable self-representation based outlier detection in a union of subspaces,” in *CVPR*, 2017, pp. 3395–3404.
- [13] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, “Estimating the support of a high-dimensional distribution,” *Neural computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [14] Y.-S. Choi, “Least squares one-class support vector machine,” *Pattern Recognition Letters*, vol. 30, no. 13, pp. 1236–1240, 2009.
- [15] T. Wang, J. Chen, Y. Zhou, and H. Snoussi, “Online least squares one-class support vector machines-based abnormal visual event detection,” *Sensors*, vol. 13, no. 12, pp. 17 130–17 155, 2013.
- [16] D. M. Tax and R. P. Duin, “Support vector data description,” *Machine learning*, vol. 54, no. 1, pp. 45–66, 2004.
- [17] D. Abati, A. Porrello, S. Calderara, and R. Cucchiara, “Latent space autoregression for novelty detection,” in *CVPR*, 2019, pp. 481–490.
- [18] R. Chalapathy, A. K. Menon, and S. Chawla, “Anomaly detection using one-class neural networks,” *arXiv preprint arXiv:1802.06360*, 2018.
- [19] W. Luo, W. Liu, and S. Gao, “A revisit of sparse coding based anomaly detection in stacked RNN framework,” in *ICCV*, 2017, pp. 341–349.
- [20] M. Hasan, J. Choi, J. Neumann, A. K. Roy-Chowdhury, and L. S. Davis, “Learning temporal regularity in video sequences,” in *CVPR*, 2016, pp. 733–742.
- [21] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, and G. Langs, “Unsupervised anomaly detection with generative adversarial networks to guide marker discovery,” in *IPMI*. Springer, 2017, pp. 146–157.
- [22] M. Ravanbakhsh, E. Sangineto, M. Nabi, and N. Sebe, “Training adversarial discriminators for cross-channel abnormal event detection in crowds,” in *WACV*. IEEE, 2019, pp. 1896–1904.
- [23] J. Wang and A. Cherian, “Learning discriminative video representations using adversarial perturbations,” in *ECCV*, 2018.
- [24] A. Edelman, T. A. Arias, and S. T. Smith, “The geometry of algorithms with orthogonality constraints,” *SIAM journal on Matrix Analysis and Applications*, vol. 20, no. 2, pp. 303–353, 1998.
- [25] W. M. Boothby, *An introduction to differentiable manifolds and Riemannian geometry*. Academic press, 1986, vol. 120.
- [26] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [27] M. M. Trivedi, T. Gandhi, and J. McCall, “Looking-in and looking-out of a vehicle: Computer-vision-based enhanced vehicle safety,” *IEEE TITS*, vol. 8, no. 1, pp. 108–120, 2007.
- [28] M. M. Trivedi, S. Y. Cheng, E. M. Childers, and S. J. Krotosky, “Occupant posture analysis with stereo and thermal infrared video: Algorithms and experimental evaluation,” *IEEE Transactions on Vehicular Technology*, vol. 53, no. 6, pp. 1698–1712, 2004.
- [29] W. Sultani, C. Chen, and M. Shah, “Real-world anomaly detection in surveillance videos,” in *CVPR*, 2018, pp. 6479–6488.
- [30] W. Li, V. Mahadevan, and N. Vasconcelos, “Anomaly detection and localization in crowded scenes,” *TPAMI*, vol. 36, no. 1, pp. 18–32, 2013.
- [31] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black, “Towards understanding action recognition,” in *ICCV*, 2013.
- [32] J. Wang and A. Cherian, “GODS: generalized one-class discriminative subspaces for anomaly detection,” in *ICCV*, 2019.
- [33] S. Matteoli, M. Diani, and J. Theiler, “An overview of background modeling for detection of targets and anomalies in hyperspectral remotely sensed imagery,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 2317–2336, 2014.
- [34] A. Lazarevic, L. Ertoz, V. Kumar, A. Ozgur, and J. Srivastava, “A comparative study of anomaly detection schemes in network intrusion detection,” in *ICDM*. SIAM, 2003.
- [35] K. A. Heller, K. M. Svore, A. D. Keromytis, and S. J. Stolfo, “One class support vector machines for detecting anomalous windows registry accesses,” in *Workshop on Data Mining for Computer Security*, 2003.
- [36] O. P. Popoola and K. Wang, “Video-based abnormal human behavior recognition: a review,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 42, no. 6, pp. 865–878, 2012.
- [37] B. Krawczyk, “One-class classifier ensemble pruning and weighting with firefly algorithm,” *Neurocomputing*, vol. 150, pp. 490–500, 2015.
- [38] B. Saleh, A. Farhadi, and A. Elgammal, “Object-centric anomaly detection by attribute-based reasoning,” in *CVPR*, 2013, pp. 787–794.
- [39] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, “A review of novelty detection,” *Signal Processing*, vol. 99, pp. 215–249, 2014.
- [40] G. Pang, C. Shen, L. Cao, and A. v. d. Hengel, “Deep learning for anomaly detection: A review,” *arXiv preprint arXiv:2007.02500*, 2020.
- [41] A. B. Tsybakov *et al.*, “On nonparametric estimation of density level sets,” *The Annals of Statistics*, vol. 25, no. 3, pp. 948–969, 1997.
- [42] D. Nolan, “The excess-mass ellipsoid,” *Journal of multivariate analysis*, vol. 39, no. 2, pp. 348–371, 1991.
- [43] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *ECCV*, 2005.

- [44] N. Dalal, B. Triggs, and C. Schmid, "Human detection using oriented histograms of flow and appearance," in *ECCV*. Springer, 2006, pp. 428–441.
- [45] D. Zhang, D. Gatica-Perez, S. Bengio, and I. McCowan, "Semi-supervised adapted hmms for unusual event detection," in *CVPR*. IEEE, 2005, pp. 611–618.
- [46] H. Xu, C. Caramanis, and S. Sanghavi, "Robust PCA via outlier pursuit," in *NIPS*, 2010, pp. 2496–2504.
- [47] J. Kim and K. Grauman, "Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates," in *CVPR*. IEEE, 2009.
- [48] M. J. Choi, A. Torralba, and A. S. Willsky, "Context models and out-of-context objects," *Pattern Recognition Letters*, vol. 33, no. 7, pp. 853–862, 2012.
- [49] S. Park, W. Kim, and K. M. Lee, "Abnormal object detection by canonical scene-based contextual model," in *ECCV*, 2012.
- [50] S. Wu, B. E. Moore, and M. Shah, "Chaotic invariants of lagrangian particle trajectories for anomaly detection in crowded scenes," in *CVPR*. IEEE, 2010, pp. 2054–2060.
- [51] S. Calderara, U. Heinemann, A. Prati, R. Cucchiara, and N. Tishby, "Detecting anomalies in peoples trajectories using spectral graph analysis," *CVIU*, vol. 115, no. 8, pp. 1099–1111, 2011.
- [52] F. Tung, J. S. Zelek, and D. A. Clausi, "Goal-based trajectory analysis for unusual behaviour detection in intelligent surveillance," *Image and Vision Computing*, vol. 29, no. 4, pp. 230–240, 2011.
- [53] L. Itti and C. Koch, "A saliency-based search mechanism for overt and covert shifts of visual attention," *Vision research*, vol. 40, no. 10–12, pp. 1489–1506, 2000.
- [54] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *ICCV*, 2009.
- [55] Y. Cong, J. Yuan, and J. Liu, "Sparse reconstruction cost for abnormal event detection," in *CVPR*. IEEE, 2011, pp. 3449–3456.
- [56] C. Lu, J. Shi, and J. Jia, "Abnormal event detection at 150 fps in Matlab," in *CVPR*, 2013, pp. 2720–2727.
- [57] B. Zhao, L. Fei-Fei, and E. P. Xing, "Online detection of unusual events in videos via dynamic sparse coding," in *CVPR*. IEEE, 2011, pp. 3313–3320.
- [58] C. Li, Z. Han, Q. Ye, and J. Jiao, "Visual abnormal behavior detection based on trajectory sparse reconstruction analysis," *Neurocomputing*, vol. 119, pp. 94–100, 2013.
- [59] H. Ren, H. Pan, S. I. Olsen, and T. B. Moeslund, "A comprehensive study of sparse codes on abnormality detection," *arXiv preprint arXiv:1603.04026*, 2016.
- [60] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *CVPR*. IEEE, 2009, pp. 248–255.
- [61] R. Girshick, "Fast R-cnn," in *ICCV*, 2015.
- [62] L. Ruff, N. Goernitz, L. Deecke, S. A. Siddiqui, R. Vandermeulen, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," in *ICML*, 2018.
- [63] D. Xu, Y. Yan, E. Ricci, and N. Sebe, "Detecting anomalous events in videos by learning deep representations of appearance and motion," *CVIU*, vol. 156, pp. 117–127, 2017.
- [64] P. Perera and V. M. Patel, "Learning deep features for one-class classification," *arXiv preprint arXiv:1801.05365*, 2018.
- [65] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," *arXiv preprint arXiv:1706.02690*, 2017.
- [66] D. Xu, E. Ricci, Y. Yan, J. Song, and N. Sebe, "Learning deep representations of appearance and motion for anomalous event detection," in *BMVC*, 2015.
- [67] Y. S. Chong and Y. H. Tay, "Abnormal event detection in videos using spatiotemporal autoencoder," in *International Symposium on Neural Networks*. Springer, 2017, pp. 189–196.
- [68] W. Luo, W. Liu, and S. Gao, "Remembering history with convolutional lstm for anomaly detection," in *ICME*. IEEE, 2017.
- [69] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *NIPS*, 2018.
- [70] S. Liang, Y. Li, and R. Srikant, "Principled detection of out-of-distribution examples in neural networks," in *ICLR*, 2018.
- [71] M. Sabokrou, M. Fayyaz, M. Fathy, and R. Klette, "Fully convolutional neural network for fast anomaly detection in crowded scenes," *arXiv preprint arXiv:1609.00866*, 2016.
- [72] M. Ravanbakhsh, M. Nabi, E. Sangineto, L. Marcenaro, C. Regazzoni, and N. Sebe, "Abnormal event detection in videos using generative adversarial nets," in *ICIP*. IEEE, 2017.
- [73] H. Park, J. Noh, and B. Ham, "Learning memory-guided normality for anomaly detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 372–14 381.
- [74] S. Mohseni, M. Pitale, J. Yadawa, and Z. Wang, "Self-supervised learning for generalizable out-of-distribution detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5216–5223.
- [75] E. Zisselman and A. Tamar, "Deep residual flow for out of distribution detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13 994–14 003.
- [76] J. Li, L. Fuxin, and S. Todorovic, "Efficient Riemannian optimization on the stiefel manifold via the cayley transform," in *ICLR*, 2020.
- [77] S. Gould, B. Fernando, A. Cherian, P. Anderson, R. S. Cruz, and E. Guo, "On differentiating parameterized argmin and argmax problems with application to bi-level optimization," *arXiv preprint arXiv:1607.05447*, 2016.
- [78] V. Lai, D. Nguyen, K. Nguyen, and T. Le, "Mixture of support vector data descriptions," in *Conference on Information and Computer Science*. IEEE, 2015.
- [79] K. Lee, D.-W. Kim, K. H. Lee, and D. Lee, "Density-induced support vector data description," *IEEE Transactions on Neural Networks*, vol. 18, no. 1, pp. 284–289, 2007.
- [80] D. M. J. Tax, "One-class classification: concept-learning in the absence of counter-examples," Ph.D. dissertation, Delft University of Technology, 2001.
- [81] F. Sohrab, J. Raitoharju, M. Gabbouj, and A. Iosifidis, "Subspace support vector data description," in *ICPR*. IEEE, 2018.
- [82] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM*, vol. 58, no. 3, p. 11, 2011.
- [83] F. De La Torre and M. J. Black, "A framework for robust subspace learning," *IJCV*, vol. 54, no. 1–3, pp. 117–142, 2003.
- [84] H. Hoffmann, "Kernel PCA for novelty detection," *Pattern recognition*, vol. 40, no. 3, pp. 863–874, 2007.
- [85] M. H. Nguyen and F. Torre, "Robust kernel principal component analysis," in *NIPS*, 2009, pp. 1185–1192.
- [86] H. Xu, C. Caramanis, and S. Mannor, "Outlier-robust PCA: the high-dimensional case," *IEEE Transactions on Information Theory*, vol. 59, no. 1, pp. 546–572, 2013.
- [87] P. Bodesheim, A. Freytag, E. Rodner, M. Kemmler, and J. Denzler, "Kernel null space methods for novelty detection," in *CVPR*, 2013, pp. 3374–3381.
- [88] W. Liu, G. Hua, and J. R. Smith, "Unsupervised one-class learning for automatic outlier removal," in *CVPR*, 2014, pp. 3826–3833.
- [89] V. Fragoso, W. Scheirer, J. Hespanha, and M. Turk, "One-class slab support vector machine," in *ICPR*. IEEE, 2016, pp. 420–425.
- [90] J. Wang, A. Cherian, F. Porikli, and S. Gould, "Video representation learning using discriminative pooling," in *CVPR*, 2018.
- [91] Y.-J. Lee and O. L. Mangasarian, "Ssvm: A smooth support vector machine for classification," *Computational optimization and Applications*, vol. 20, no. 1, pp. 5–22, 2001.
- [92] O. L. Mangasarian and D. R. Musicant, "Lagrangian support vector machines," *Journal of Machine Learning Research*, vol. 1, no. Mar, pp. 161–177, 2001.
- [93] H. Jégou, M. Douze, and C. Schmid, "On the burstiness of visual elements," in *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 2009, pp. 1169–1176.
- [94] P. Koniusz, F. Yan, P.-H. Gosselin, and K. Mikolajczyk, "Higher-order occurrence pooling for bags-of-words: Visual concept detection," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 2, pp. 313–326, 2016.
- [95] A. B. Graf, A. J. Smola, and S. Borer, "Classification in a normalized feature space using support vector machines," *IEEE Transactions on Neural Networks*, vol. 14, no. 3, pp. 597–605, 2003.
- [96] R. Ranjan, C. D. Castillo, and R. Chellappa, "L2-constrained softmax loss for discriminative face verification," *arXiv preprint arXiv:1703.09507*, 2017.
- [97] A. Cherian, S. Sra, S. Gould, and R. Hartley, "Non-linear temporal subspace representations for activity recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2197–2206.
- [98] D. Tax and R. Duin, "Feature scaling in support vector data descriptions," *Learning from Imbalanced Datasets*, pp. 25–30, 2000.
- [99] R. Herbrich and T. Graepel, "A pac-bayesian margin bound for linear classifiers: Why svms work," *Advances in neural information processing systems*, pp. 224–230, 2001.
- [100] A. J. Smola and B. Schölkopf, *Learning with kernels*. Citeseer, 1998, vol. 4.

- [101] P.-A. Absil and K. A. Gallivan, "Joint diagonalization on the oblique manifold for independent component analysis," in *ICASSP*. IEEE, 2006.
- [102] Z. Wen and W. Yin, "A feasible method for optimization with orthogonality constraints," *Mathematical Programming*, vol. 142, no. 1-2, pp. 397–434, 2013.
- [103] R. J. Muirhead, *Aspects of multivariate statistical theory*. John Wiley & Sons, 2009, vol. 197.
- [104] N. J. Higham, C. Mehl, and F. Tisseur, "The canonical generalized polar decomposition," *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 4, pp. 2163–2180, 2010.
- [105] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a matlab toolbox for optimization on manifolds," *JMLR*, vol. 15, no. 1, pp. 1455–1459, 2014.
- [106] L. S. Nordhoff, *Motor vehicle collision injuries: biomechanics, diagnosis, and management*. Jones & Bartlett Learning, 2005.
- [107] S. M. Duma, T. A. Kress, D. J. Porta, C. D. Woods, J. N. Snider, P. M. Fuller, and R. J. Simmons, "Airbag-induced eye injuries: a report of 25 cases," *Journal of Trauma and Acute Care Surgery*, vol. 41, no. 1, pp. 114–119, 1996.
- [108] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, "Realtime multi-person 2d pose estimation using part affinity fields," *arXiv preprint arXiv:1611.08050*, 2016.
- [109] T. S. Kim and A. Reiter, "Interpretable 3d human action analysis with temporal convolutional networks," in *CVPRW*, 2017.
- [110] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "NTU RGB+ D: A large scale dataset for 3d human activity analysis," in *CVPR*, 2016.
- [111] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *CVPR*. IEEE, 2017, pp. 4724–4733.
- [112] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [113] J. Townsend, N. Koep, and S. Weichwald, "Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation," *JMLR*, vol. 17, no. 137, pp. 1–5, 2016.
- [114] A. Moh'd A Mesleh, "Chi square feature extraction based svms arabic language text categorization system," *Journal of Computer Science*, vol. 3, no. 6, pp. 430–435, 2007.
- [115] A. Barla, F. Odone, and A. Verri, "Histogram intersection kernel for image classification," in *ICIP*, vol. 3. IEEE, 2003, pp. III–513.
- [116] M. Ravanbakhsh, M. Nabi, H. Mousavi, E. Sangineto, and N. Sebe, "Plug-and-play cnn for crowd motion analysis: An application in abnormal event detection," in *Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1689–1698.
- [117] V. Hautamaki, I. Karkkainen, and P. Franti, "Outlier detection using k-nearest neighbour graph," in *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, vol. 3. IEEE, 2004, pp. 430–433.



several awards, including the Best Student Paper Award at ICIP, 2012.

Anoop Cherian is a Principal Research Scientist with Mitsubishi Electric Research Labs (MERL) Cambridge, MA and an Adjunct Researcher with the Australian National University. He received his M.S. and Ph.D. degrees in computer science from the University of Minnesota, Minneapolis in 2010 and 2013, respectively. He has broad interests in machine learning, computer vision, deep learning, and non-Euclidean optimization. He has co-authored more than 50 scientific articles, and is also the recipient of



Jue Wang is a PhD student with the Research School of Engineering at the Australian National University since 2016. He is also associated with CSIRO's Data61 in Australia. From 2010-2014, he received his double bachelor degree (honors) in Electronic Engineering from Australian National University and Beijing Institute of Technology. His research interest are in the area of computer vision and machine learning.