



DOC³: deep one class classification using contradictions

Sauptik Dhar¹ · Bernardo Gonzalez-Torres²

Received: 27 June 2022 / Revised: 12 April 2023 / Accepted: 3 July 2023 /
Published online: 29 August 2023

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2023

Abstract

This paper introduces the notion of learning from contradictions (a.k.a. Universum learning) for deep one class classification problems. We formalize this notion for the widely adopted one class large-margin loss (Schölkopf et al. in *Neural Comput* 13(7):1443–1471), and propose the deep one class classification using contradictions (DOC³) algorithm. We show that learning from contradictions incurs lower generalization error by comparing the empirical Rademacher complexity of DOC³ against its traditional inductive learning counterpart. Further, our proposed ‘learning from contradiction’ is a generic learning setting and can compliment other advanced learning settings. To illustrate this, we extend the adversarial learning based DROCC-LF (Goyal et al. in *International conference on machine learning, PMLR, 2020*) algorithm under this new setting. Our empirical results demonstrate the efficacy of DOC³ and its extensions compared to popular baseline algorithms on several benchmark and real-life data sets.

Keywords Anomaly detection · One class classification · Universum learning

1 Introduction

Anomaly detection (AD) is one of the most widely researched problem in the machine learning community (Chandola et al., 2009). In its basic form, the task of Anomaly Detection (AD) involves discerning patterns in data that do not conform to expected ‘normal’ behavior. These non-conforming patterns are referred to as anomalies or outliers. Anomaly detection problems manifest in several forms in real-life like, defect detection in manufacturing lines, intrusion detection for cyber security, or pathology detection for medical diagnosis etc. There are several mechanisms to handle anomaly detection problems viz., parametric or non-parametric statistical modeling, spectral based, or classification based modeling (Chan-

✉ Sauptik Dhar
sauptik.dhar@gmail.com

Bernardo Gonzalez-Torres
bernardo.gtorres@gmail.com

¹ Eluvio, Berkeley, CA, USA

² CIC, Instituto Politécnico Nacional, Mexico City, Mexico

dola et al., 2009). Of these, the classification based approach has been widely adopted in literature (Scholkopf et al., 2002; Tax & Duin, 2004; Tan et al., 2016; Cherkassky & Mulier, 2007). One specific classification based formulation which has gained huge adoption is one class classification (Scholkopf et al., 2002; Tax & Duin, 2004), where we design a parametric model to estimate the support of the ‘normal’ class distribution. The estimated model is then used to detect ‘unseen’ abnormal samples.

With the recent success of deep learning based approaches for different machine learning problems, there has been a surge in research adopting deep learning for one class problems (Ruff et al., 2021; Pang et al., 2020; Chalapathy & Chawla, 2019). However, most of these works adopt an inductive learning setting. This makes the underlying model estimation data hungry, and perform poorly for applications with limited training data availability, like medical diagnosis, industrial defect detection, etc. The *learning from contradictions* paradigm (popularly known as Universum learning) has shown to be particularly effective for problems with limited training data availability (Vapnik, 2006; Sinz et al., 2008; Weston et al., 2006; Chen & Zhang, 2009; Cherkassky et al., 2011; Shen et al., 2012; Dhar & Cherkassky, 2015; Zhang & LeCun, 2017; Xiao et al., 2021). However, it has been mostly limited to binary or multi class problems. In this paradigm, along with the labeled training data we are also given a set of unlabeled contradictory (a.k.a universum) samples. These universum samples belong to the same application domain as the training data, but are known not to belong to any of the classes. The rationale behind this setting comes from the fact that even though obtaining labels is very difficult, obtaining such additional unlabeled samples is relatively easier. These unlabeled universum samples act as *contradictions* and should not be explained by the estimated decision rule. Adopting this to one class problems is not straight forward. A major conceptual problem is that, one class model estimation represents unsupervised learning, where the notion of contradiction needs to be redefined properly. In this paper,

1. *Definition* We introduce the notion of ‘Learning from contradictions’ for one class problems (Definition 2).
2. *Formulation* We analyze the popular one class hinge loss (Schölkopf et al., 2001), and extend it under universum settings to propose the Deep One Class Classification using Contradictions DOC³ algorithm. Further, our proposed ‘learning from contradiction’ is a generic learning setting and can compliment other advanced learning settings. To illustrate this, we extend the adversarial learning based DROCC-LF (Goyal et al., 2020) algorithm under universum settings and call it DROCC-LF (univ) (see Algo. 1).
3. *Generalization error* We analyze the generalization performance of one class formulations under inductive and universum settings using Rademacher complexity based bounds, and show that learning under the universum setting can provide improved generalization compared to its inductive counterpart.
4. *Empirical results* Finally, we provide an exhaustive set of empirical results on several tabular and image datasets in support of our approach.

2 One class learning under inductive settings

First we introduce the widely adopted inductive learning setting used for one class problems (Scholkopf et al., 2002; Cherkassky & Mulier, 2007).

Definition 1 (*Inductive setting*) Given i.i.d training samples from a single class $\mathcal{T} = (\mathbf{x}_i, y_i = +1)_{i=1}^n \sim \mathcal{D}_{\mathcal{X}|\mathcal{Y}=+1}^n$, with $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ and $y \in \mathcal{Y} = \{-1, +1\}$; estimate a

hypothesis $h^* : \mathcal{X} \rightarrow \mathcal{Y}$ from an hypothesis class \mathcal{H} which minimizes,

$$\inf_{h \in \mathcal{H}} \mathbb{E}_{\mathcal{D}_T} [\mathbf{1}_{y \neq h(\mathbf{x})}] \tag{1}$$

\mathcal{D}_T is the training distribution (consisting of both classes)

$\mathcal{D}_{\mathcal{X}|\mathcal{Y}=+1}$ is class conditional distribution

$\mathbf{1}(\cdot)$ is the indicator function, and

$\mathbb{E}_{\mathcal{D}_T}(\cdot)$ is the expectation under training distribution.

Note that, the underlying data generation process assumes a two class problem; of which the samples from only one class is available during training. The overall goal is to estimate a model which minimizes the error on the future test data, containing samples from both normal ($y = +1$) and abnormal classes ($y = -1$). Typical examples include, AI driven visual inspection of product defects in a manufacturing line; where images or videos of non-defective products are available in abundance. The goal is to detect ‘defective’ (abnormal/anomalous) products through visual inspection in manufacturing lines (Bergmann et al., 2019; Weimer et al., 2016). A popular loss function used in such settings is the ν -SVM loss (Schölkopf et al., 2001),

$$\begin{aligned} \min_{\mathbf{w}, \xi, \rho} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + \frac{1}{\nu n} \sum_{i=1}^n \xi_i - \rho \\ \text{s.t.} \quad & \mathbf{w}^\top \phi(\mathbf{x}_i) \geq \rho - \xi_i, \quad \xi_i \geq 0; \quad \forall i = 1 \dots n \end{aligned} \tag{2}$$

where $\nu \in (0, 1]$ is a user-defined parameter which controls the margin errors $\sum_i \xi_i$ and the size of geometric $\frac{1}{\|\mathbf{w}\|}$ and functional ρ margins. $\phi(\cdot) : \mathcal{X} \rightarrow \mathcal{G}$ is a feature map. Typical examples include an *empirical* kernel map (see Definition 2.15 (Scholkopf et al., 2002)) or a map induced by a deep learning network (Goodfellow et al., 2016). The final decision function is given as, $h(\mathbf{x}) = \begin{cases} +1; & \text{if } \mathbf{w}^\top \phi(\mathbf{x}_i) \geq \rho \\ -1; & \text{else} \end{cases}$. Note that, recent works like Ruff et al. (2018) extend a different loss function which uses a ball to explain the support of the data distribution following (Tax & Duin, 2004). As discussed in Schölkopf et al. (2001), most of the time these two formulations yield equivalent decision functions. For example, with kernel machines $\mathbf{K}(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^\top \phi(\mathbf{x}')$ depending solely on $\mathbf{x} - \mathbf{x}'$ (like RBF kernels), these two formulations are the same. Hence, most of the improvements discussed in this work translates to such alternate formulations. In this paper however, we solve the following one class Hinge Loss,

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + C L_T(\mathbf{w}, \{\phi(\mathbf{x}_i)\}_{i=1}^n) \\ \text{s.t} \quad & L_T(\mathbf{w}, \{\phi(\mathbf{x}_i)\}_{i=1}^n) = \sum_{i=1}^n [1 - \mathbf{w}^\top \phi(\mathbf{x}_i)]_+; \quad [x]_+ = \max(0, x) \end{aligned} \tag{3}$$

to estimate the the decision function $f(\mathbf{x}) = \mathbf{w}^\top \phi(\mathbf{x}_i)$ and use the decision rule, $h(\mathbf{x}) = \begin{cases} +1; & \text{if } f(\mathbf{x}) \geq 1 \\ -1; & \text{else} \end{cases}$. Here, the user-defined parameter C controls the trade-off between explaining the training samples (through small margin error $\sum_{i=1}^n \xi_i$), and the margin size (through $\|\mathbf{w}\|_2^2$), which in turn controls the generalization error. For deep learning architectures we optimize using all the model parameters and equivalently regularize the entire matrix norm $\|\mathbf{W}\|_F^2$, see Goyal et al. (2020), Ruff et al. (2018). Note that, we solve one class Hinge loss (3) for the two main reasons,

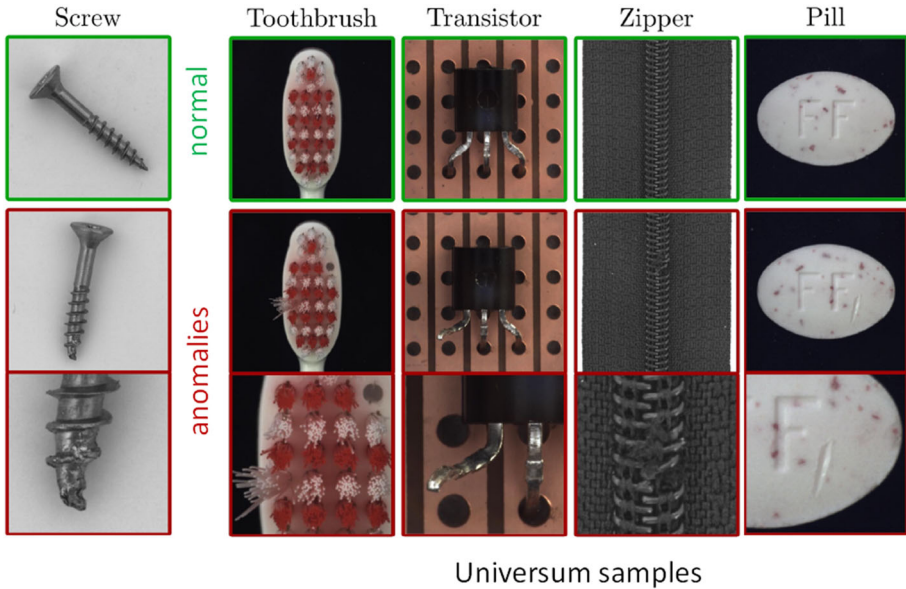


Fig. 1 Visual inspection of anomalous screws in a manufacturing line (Bergmann et al., 2019). Images of the other products act as universum samples. Such images are neither normal-screw nor anomalous-screw images and act as contradictions

- First, it has the advantage that $L_T(\mathbf{w}, \phi(\{\mathbf{x}\}_{i=1}^n)) = \sum_{i=1}^n [1 - \mathbf{w}^T \phi(\mathbf{x}_i)]_+$ exhibits the same form as the traditional hinge loss used for binary classification problems (Vapnik, 2006) and can be easily solved using existing software packages (Paszke et al., 2019; Abadi et al., 2016; Pedregosa et al., 2011). Throughout the paper we refer (3) using underlying deep architectures as Deep One Class DOC (Hinge) formulation.
- Second, solving Eq. (3) also provides the solution for Eq. (2). This connection follows from Proposition 1.

Proposition 1 Connection between Eq. (2) and Eq. (3)

1. Any solution \mathbf{w} of Eq. (3) also solves Eq. (2) with $\nu = \frac{1}{Cn\delta}$; where $\delta > 0$ is a scalar that depends on the solution of Eq. (3). Further, this solution $(\hat{\mathbf{w}}, \rho)$ of Eq. (2) is given as $\hat{\mathbf{w}} = \mathbf{w}\delta, \rho = \delta$.
2. The decision function obtained through solving Eq. (3) i.e., $\mathbf{w}^T \phi(\mathbf{x}) - 1 = 0$ coincides with the decision function $\hat{\mathbf{w}}^T \phi(\mathbf{x}) - \rho = 0$ obtained by solving Eq. (2) i.e. using the solution discussed above.

All proofs are provided in “Appendix”.

3 One class learning using contradictions a.k.a Universum learning

3.1 Problem formulation

Learning from contradictions or Universum learning was introduced in Vapnik (2006) for binary classification problems to incorporate a priori knowledge about admissible data

samples. For example, if the goal of learning is to discriminate between handwritten digits ‘5’ and ‘8’, one can introduce additional knowledge in the form of other handwritten letters ‘a’, ‘b’, ‘c’, ‘d’, . . . ‘z’. These examples from the Universum contain certain information about the handwritten styles of authors, but they cannot be assigned to any of the two classes (5 or 8). Further, these Universum samples do not have the same distribution as labeled training samples. In this work we introduce the notion of ‘Learning from Contradictions’ for one class problems. Similar to inductive setting (Definition 1) the goal here is also to minimize the generalization error on future test data containing both normal ($y = +1$) and abnormal ($y = -1$) samples. Here however, during training in addition to the samples from the normal class $(\mathbf{x}_i, y_i = +1)_{i=1}^n$, we are also provided with universum (contradictory) samples, which are known not to belong to either of the (normal or abnormal) classes of interest. A practical use-case can be of automated visual inspection based anomaly detection in manufacturing lines. Here the target is to identify the defects in a specific product type (say ‘screws’ in Fig. 1). For this case, the images from other product types in the manufacturing line act as universum samples. Note that, such universum samples belong to the same application domain (i.e. visual inspection data); but do not represent either of the classes normal screws or anomalous screws. This setting is formalized as,

Definition 2 (*Learning from contradictions a.k.a Universum setting*) Given i.i.d training samples $\mathcal{T} = (\mathbf{x}_i, y_i = +1)_{i=1}^n \sim \mathcal{D}_{\mathcal{X}|\mathcal{Y}=+1}^n$, with $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ and $y \in \mathcal{Y} = \{-1, +1\}$ and additional m universum samples $\mathcal{U} = (\mathbf{x}_{i'}^*)_{i'=1}^m \sim \mathcal{D}_{\mathcal{U}}$ with $\mathbf{x}^* \in \mathcal{X}_{\mathcal{U}}^* \subseteq \mathbb{R}^d$, estimate $h^* : \mathcal{X} \rightarrow \mathcal{Y}$ from hypothesis class \mathcal{H} which, in addition to Eq. (1), obtains maximum contradiction on universum samples i.e. maximizes the following probability for $\mathbf{x}^* \in \mathcal{X}_{\mathcal{U}}^*$,

$$\sup_{h \in \mathcal{H}} \mathbb{P}_{\mathcal{D}_{\mathcal{U}}} [h(\mathbf{x}^*) \notin \mathcal{Y}] = \sup_{h \in \mathcal{H}} \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} [\mathbf{1}_{\{\bigcap_{y \in \mathcal{Y}} h(\mathbf{x}^*) \neq y\}}] \tag{4}$$

$\mathcal{D}_{\mathcal{U}}$ is the universum distribution,

$\mathbb{P}_{\mathcal{D}_{\mathcal{U}}}(\cdot)$ is probability under universum distribution,

$\mathbb{E}_{\mathcal{D}_{\mathcal{U}}}(\cdot)$ is the expectation under universum distribution, $\mathcal{X}_{\mathcal{U}}^*$ is the domain of universum data.

Learning using contradictions under Universum setting has the dual goal of minimizing the generalization error in Eq. (1) while maximizing the contradiction on universum samples Eq. (4). The following proposition provides guidelines on how this can be achieved for the one class hinge loss in Eq. (3).

Proposition 2 *For the one class hinge loss in Eq. (3), maximum contradiction on universum samples $\mathbf{x}^* \in \mathcal{X}_{\mathcal{U}}^*$ can be achieved when,*

$$|\mathbf{w}^\top \phi(\mathbf{x}^*) - 1| = 0 \tag{5}$$

That is, we need the universum samples to lie on the decision boundary. This motivates the following one class loss using contradictions (under Universum settings) where we relax the constraint in Eq. (5) by introducing a Δ -insensitive loss similar to Weston et al. (2006), Dhar et al. (2019) and solve,

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \|\mathbf{w}\|_2^2 + C L_T(\mathbf{w}, \phi(\{\mathbf{x}_i\}_{i=1}^n)) + C_U L_U(\mathbf{w}, \phi(\{\mathbf{x}_{i'}^*\}_{i'=1}^m)) \\ \text{s.t.} \quad & L_T(\mathbf{w}, \phi(\{\mathbf{x}_i\}_{i=1}^n)) = \sum_{i=1}^n [1 - \mathbf{w}^\top \phi(\mathbf{x}_i)]_+ \\ & L_U(\mathbf{w}, \phi(\{\mathbf{x}_{i'}^*\}_{i'=1}^m)) = \sum_{i'=1}^m [|1 - \mathbf{w}^\top \phi(\mathbf{x}_{i'}^*)| - \Delta]_+ \end{aligned} \tag{6}$$

Here, $[x]_+ = \max(0, x)$. Further, the interplay between C, C_U controls the trade-off between explaining the training samples using L_T versus maximizing the contradiction on Universum samples using L_U . For $C_U = 0$ or $\Delta \rightarrow \infty$, Eq. (6) transforms to Eq. (3). For deep learning models, we optimize Eq. (6) over all the model parameters and refer to it as Deep One Class Classification using Contradictions (DOC³).

3.2 Analysis of generalization error bound

Next we provide theoretical justification in support of Universum learning. We argue, learning under universum settings using DOC³ can provide improved generalization error compared to its inductive counterpart DOC (Hinge). For this, we first derive a generic form of the generalization error bound for one class learning using the Rademacher complexity capacity measure in Theorem 1.

Theorem 1 (Generalization error bound) *Let \mathcal{F} be the class of functions from which the decision function $f(\mathbf{x})$ in Eq. (3) and (6) are estimated. Let $R_{f,1} = \{\mathbf{x} : f(\mathbf{x}) \geq 1\}$ be the induced decision region. Then, with probability $1 - \eta$ with $\eta \in [0, 1]$, over any independent draw of the random sample $\mathcal{T} = (\mathbf{x}_i, y_i = +1)_{i=1}^n \sim \mathcal{D}_{\mathcal{T}|y=+1}^n$, for any $\kappa > 0$ we have,*

$$\mathbb{P}_{\mathcal{D}_{\mathcal{T}|y=+1}}(\mathbf{x} \notin R_{f,1-\kappa}) \leq \frac{1}{\kappa n} \sum_{i=1}^n \xi_i + \frac{2}{\kappa} \hat{\mathcal{R}}_n(\mathcal{F}) + 3\sqrt{\frac{\ln \frac{2}{\eta}}{2n}} \tag{7}$$

where $\xi_i = [1 - f(\mathbf{x}_i)]_+$; $R_{f,\theta} = \{\mathbf{x} : f(\mathbf{x}) \geq \theta\}$

$$\hat{\mathcal{R}}_n(\mathcal{F}) = \mathbb{E}_{\sigma} \left[\sup_{f \in \mathcal{F}} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right| \middle| (\mathbf{x}_i)_{i=1}^n \right]$$

$\sigma =$ independent uniform $\{\pm 1\}$ - valued random variables a.k.a Rademacher variables.

The Theorem 1 is agnostic of model parameterization and holds for any popularly adopted kernel machine or deep learning architectures. Similar to the Theorem 7 in Schölkopf et al. (2001), Theorem 1 gives a probabilistic guarantee that new points lie in a larger region $R_{f,1-\kappa}$. Here, we rather use the Empirical Rademacher Complexity (ERC) $\hat{\mathcal{R}}_n(\mathcal{F})$ as the capacity measure of the hypothesis class, instead of the covering number. Additionally, our bound does not contain a $\frac{1}{\kappa^2}$ term as in Schölkopf et al. (2001), and only has the scaling factor of $\frac{1}{\kappa}$. As seen from Theorem 1 above, it is preferable to use a hypothesis class \mathcal{F} with smaller ERC $\hat{\mathcal{R}}_n(\mathcal{F})$. Next we compare the ERC of the hypothesis class induced by the formulations Eq. (3) versus Eq. (6).

Theorem 2 (Empirical Rademacher complexity). *For the hypothesis class induced by the formulations,*

- Equation (3): $\mathcal{F}_{ind} = \{f : \mathbf{x} \rightarrow \mathbf{w}^\top \phi(\mathbf{x}) \mid \|\mathbf{w}\|_2^2 \leq \Lambda^2\}$
- Equation (6): $\mathcal{F}_{univ} = \{f : \mathbf{x} \rightarrow \mathbf{w}^\top \phi(\mathbf{x}) \mid \|\mathbf{w}\|_2^2 \leq \Lambda^2; |\mathbf{w}^\top \phi(\mathbf{x}^*) - 1| \leq \Delta, \forall \mathbf{x}^* \in \mathcal{X}_U^*\}$

The following holds,

1. $\hat{\mathcal{R}}_n(\mathcal{F}_{ind}) \geq \hat{\mathcal{R}}_n(\mathcal{F}_{univ})$
2. Further, for any fixed mapping $\phi(\cdot)$, $\forall \gamma \geq 0$ we have,

$$(a) \hat{\mathcal{R}}_n(\mathcal{F}_{ind}) \leq \frac{2\Delta}{n} \sqrt{\sum_{i=1}^n \|\mathbf{z}_i\|^2}; \text{ where } \mathbf{z} = \phi(\mathbf{x})$$

$$(b) \hat{\mathcal{R}}_n(\mathcal{F}_{univ}) \leq \frac{2\Delta}{n} \sqrt{\sum_{i=1}^n \|\mathbf{z}_i\|^2} \min_{\gamma \geq 0} K(\gamma) [1 - \Sigma(\gamma)]^{\frac{1}{2}}$$

$$\text{where } K(\gamma) = \left[1 + \frac{2\gamma m(\Delta^2 + 1)}{\Delta^2}\right]^{\frac{1}{2}} \quad (8)$$

$$\Sigma(\gamma) = \gamma \frac{\text{tr}(VZ^T ZV^T)}{[\text{tr}(Z^T Z)] [\text{tr}(I + \gamma VV^T)]}$$

$$Z = \begin{bmatrix} (\mathbf{z}_1)^T \\ \vdots \\ (\mathbf{z}_n)^T \end{bmatrix} \text{ and } V = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \begin{bmatrix} (\mathbf{u}_1)^T \\ \vdots \\ (\mathbf{u}_m)^T \end{bmatrix}; \quad \mathbf{u} = \phi(\mathbf{x}^*); \quad \mathbf{x}^* \in \mathcal{X}_U^* \quad (9)$$

$\otimes =$ Kronecker Product, $\text{tr} =$ Matrix Trace

Note that, several recent works (Neyshabur et al., 2015; Sokolic et al., 2016; Cortes et al., 2017) derive the ERC of the function class induced by an underlying neural architecture. In this analysis however, we fix the feature map and analyze how the loss function in Eq. (6) reduces the function class capacity compared to Eq. (3). This simplifies our analysis and focuses on the effect of the proposed new loss in Eq. (6) under the universum setting. As seen from Theorem 2 (1), the function class induced under the universum setting (using contradictions) exhibits lower ERC compared to that under inductive settings. A more explicit characterization of the ERC is provided in part (2). Setting $\gamma = 0$ in (b), we achieve the same R.H.S as (a); hence the R.H.S in (b) is always smaller than in (a). Further note that $\Sigma(\gamma)$ in Eq. (9) has the form of a correlation matrix between the training and universum samples in the feature space. In fact, we have $\Sigma(\infty) = \lim_{\gamma \rightarrow \infty} \Sigma(\gamma) = \frac{\text{tr}(VZ^T ZV^T)}{\text{tr}(Z^T Z) \text{tr}(VV^T)}$. This shows that, for a fixed number of universum samples m and Δ , the effect of the DOC³ algorithm is influenced by the correlation between training and universum samples in the feature space. Loosely speaking, the DOC³ algorithm searches for a solution where in addition to reducing the margin errors ξ_i , also minimizes this correlation; and by doing so minimizes the generalization error. Similar conclusions have been empirically derived for binary, multiclass problems in Weston et al. (2006), Chapelle et al. (2008), Cherkassky et al. (2011) and Dhar et al. (2019). Here, we provide the theoretical reasoning for one class problems. Further, we confirm these theoretical findings in our results (Sect. 5.3.3).

3.3 Algorithm implementation

A limitation in solving Eq. (6) is handling the absolute term in L_U . In this paper we adopt a similar approach used in Weston et al. (2006), Dhar et al. (2019) and simplify this by re-writing L_U as a sum of two hinge functions. To do this, for every universum sample \mathbf{x}_i^* we create two artificial samples, $(\mathbf{x}_{i1}^*, y_{i1}^* = 1)$, $(\mathbf{x}_{i2}^*, y_{i2}^* = -1)$ and re-write,

$$\begin{aligned} L_U &= \sum_{i'=1}^m [|1 - \mathbf{w}^T \phi(\mathbf{x}_{i'}^*)| - \Delta]_+ \\ &= \sum_{i'=1}^m \left([\epsilon_1 - y_{i'1}^* \mathbf{w}^T \phi(\mathbf{x}_{i'}^*)]_+ + [\epsilon_2 - y_{i'2}^* \mathbf{w}^T \phi(\mathbf{x}_{i'}^*)]_+ \right) \end{aligned} \quad (10)$$

where, $\epsilon_1 = 1 - \Delta$ and $\epsilon_2 = -1 - \Delta$. Now, the universum loss is the sum of two hinge functions with ϵ_1, ϵ_2 - margins; and can be solved using standard deep learning libraries (Paszke et al., 2019; Abadi et al., 2016; Pedregosa et al., 2011).¹

4 Existing approaches and related works

Most research in Anomaly Detection (AD) can be broadly categorized as adopting either traditional (shallow) or the more modern deep learning based approaches. Traditional approaches generally adopt parametric or non-parametric statistical modeling, spectral based, or classification based modeling (Chandola et al., 2009). Typical examples include, PCA based methods (Jolliffe, 2002; Hoffmann, 2007), proximity based methods (Knorr et al., 2000; Ramaswamy et al., 2000), tree-based methods like Isolation Forest (IF) (Liu et al., 2008), or classification based OC-SVM (Schölkopf et al., 2001), Support Vector Data Description (SVDD) (Tax & Duin, 2004) etc. These techniques provide good performance for optimally tuned feature map. However, for complex domains like vision or speech, where designing optimal feature maps is non trivial; such approaches perform sub-optimally. A detailed survey on these approaches is available in Chandola et al. (2009).

In contrast, for the modern deep learning based approaches, extracting the optimal feature map is imbibed in the learning process. Broadly there are three main sub-categories for deep learning based AD. First, the Deep Auto Encoder and its variants like DCAE (Masci et al., 2011; Makhzani Frey, 2014) or ITAE (Huang et al., 2019) etc. Here, the aim is to build an embedding where the normal samples are correctly reconstructed while the anomalous samples exhibit high reconstruction error. The second type of approach adopt Generative Adversarial Network (GAN)-based techniques like AnoGAN (Schlegl et al., 2017), GANomaly (Akçay et al., 2018), EGBAD (Zenati et al., 2018), CBiGAN (Carrara et al., 2020) etc. These approaches, typically focus on generating additional samples which follow similar distribution as the training data. This is followed up by designing an anomaly score to discriminate between normal versus anomalous samples. Finally, the third category consist of the more recent one class classification based approaches like, Deep SVDD (Ruff et al., 2018), DROCC (Goyal et al., 2020) etc. These approaches adopt solving a one class loss function catered for deep architectures. All these above approaches however adopt an unsupervised inductive learning setting. There is a newer class of classification based paradigm which adopts semi or self supervised formulations. Typical examples include, GOAD (Bergman Hoshen, 2020), SSAD (Ruff et al., 2019), ESAD (Huang et al., 2020) etc. However, such approaches use fundamentally different problem settings (like a multi class problem for GOAD); or have different assumptions on the additional data available.

Learning with disjoint auxiliary (DA) data: A recently popularized new learning setting assumes the availability of an additional auxiliary data which is disjoint from the test set. The underlying assumption is that these auxiliary samples may or may not follow the same distribution as the test data and are disjoint from test set. This idea was first introduced in Dhar (2014) (see Sect. 4.3) and misconstrued as Universum learning. Note that, the notion of universum samples was originally introduced to act as contradictions to the concept classes in the test set (Vapnik, 2006). The above assumption does not adhere to this notion and violates the true essence of Universum learning. This setting has been recently used to propose ‘outlier exposure’ in Hendrycks et al. (2018) and variants (Ruff et al., 2021; Goyal et al., 2020). A more advanced variation of this setup adopts generating the anomalous samples through

¹ All codes are available at: <https://github.com/sauptikdhar/DOC3>

perturbation (Cai & Fan, 2022) or through distribution-shifting transformations (Tack et al., 2020) and using contrastive losses. Our learning from contradiction setting is different from the above methods in the following aspects,

- (*Problem setting*) is different. While the above setting only assumes disjoint auxiliary data from test data's concept classes ('normal' and 'anomalous' samples), Universum follows a different assumption that the concept classes of the universum data is different from both the normal as well as anomalous samples. This assumption is quintessential for proving Prop. 2, which in turn provides the optimality constraint on the decision function (in Eq. 5). Prop. 2 is not possible for DA setting.
- (*Formulation*) The difference in problem setting is also clear from the formulations. For example, the formulations proposed under the disjoint auxiliary setting like, (Dhar, 2014), Outlier Exposure (OE) (Hendrycks et al., 2018), DROCC-LF (OE) (Goyal et al., 2020) etc., only uses the relation between in-lie training data and the additional auxiliary data. No information on the relation between the auxiliary data and the anomalous samples in test set is encoded in the loss function. In essence, such approaches controls the complexity of hypotheses class by constraining the space in which 'normal' samples can lie. In contrast, Universum learning assumes different concept classes for Universum versus both normal and anomalous (test) samples. This information is encoded through the proof in Prop. 2. The Universum setting controls the complexity of hypotheses class by constraining the space in which both 'normal' or 'anomalous' samples can lie.

In short, Universum learning adopts a different learning paradigm (see Definition 2) compared to the 'disjoint auxiliary data' settings. Different from the existing 'disjoint auxiliary' based loss functions in Dhar (2014), OE (Hendrycks et al., 2018), DROCC-LF (OE) (Goyal et al., 2020) etc., the Universum samples (in Eq. (6)) implicitly contradicts the unseen anomalous test samples. A more pedagogical explanation of the differences between these settings with examples is provided in "Appendix C.1". However similar as DA/OE settings, Universum learning can complement other advanced learning settings. To highlight this, we extend the adversarial based DROCC-LF algorithm under universum setting in algorithm 1 and compare it's performance against its OE based extension DROCC-LF(OE) (introduced in Goyal et al. (2020)). Here, for DROCC-LF (univ) we replace the binary cross entropy loss used in Goyal et al. (2020) with the universum loss in Eq. (6) (see step 3 in Algo. 1). We use the same notations used in Goyal et al. (2020).

5 Empirical results

5.1 Standard benchmark on tabular datasets from Goyal et al. (2020)

First we provide the results on several tabular data used in Goyal et al. (2020). The datasets used involves standard anomaly detection problems described below,

- *Abalone* used in Das et al. (2018): Here the task is to predict the age of abalone using several physical measurements like, rings, sex, length, diameter, height, weight, etc. For this problem class 3 and 21 are anomalies and class 8, 9 and 10 serve as normal samples.
- *Arrhythmia* used in Zong et al. (2018). Here the task is to identify the arrhythmic samples using the ECG features. We follow the same data set preparation as Zong et al. (2018).
- *Thyroid* used in Zong et al. (2018). The goal is to predict if a patient is hypothyroid based on his/her medical history. We follow the same data set preparation as Zong et al. (2018).

Algorithm 1 DROCC-LF (univ)

Input: Training (normal) samples $\mathcal{T} = (\mathbf{x}_i, y_i = +1)_{i=1}^n$ and Universum samples $\mathcal{U} = (\mathbf{x}_{i'}^*)_{i'=1}^m$.
Parameters: Radius $r, \lambda \geq 0, \mu \geq 0$, step-size η , number of gradient steps m_g , number of initial training steps n_0 .
Initial steps: For $B = 1, \dots, n_0$
 Batch of training (X_T) and universum (X_U) samples

$$\theta = \theta - \nabla \left(\sum_{\mathbf{x}_i \in X_B} L_T(f(\mathbf{x}_i)) + \sum_{\mathbf{x}_{i'}^* \in X_U} L_U(f(\mathbf{x}_{i'}^*)) \right)$$

DROCC steps: For $B = n_0, \dots, n_0 + N$
 X_T : Batch of normal training inputs ($y = +1$)
 $\forall x \in X_T : h \sim \mathcal{N}(0, I_d)$
Adversarial search: For $i = 1, \dots, m_g$
 1. $L_T(h) = L_T(f(x + h), -1)$
 2. $h = h + \eta \frac{\nabla_h L_T(h)}{\|\nabla_h L_T(h)\|}$
 3. $h =$ Projection given by Prop.1 in Goyal et al. (2020)

$$e^{itr} = \lambda \|\mathbf{w}\|^2 + \sum_{\mathbf{x}_i \in X_B} L_T(f(\mathbf{x}_i)) + \sum_{\mathbf{x}_{i'}^* \in X_U} L_U(f(\mathbf{x}_{i'}^*)) + \mu L_T(f(x + h), -1)$$

$$\theta = \theta - \nabla e^{itr}$$

For all the above data we use the data set preparation codes provided in Goyal et al. (2020). This code provides the data preprocessing and partitioning scheme as used in the previous works. We follow the same experiment setup and network architecture as in Goyal et al. (2020). We use the same baseline methods as used in Goyal et al. (2020), and also provide the results of the recent approach PLAD (Cai & Fan, 2022) proposed to stabilize the DROCC baseline.

Table 1 provides the results of DOC³ over 10 random partition of the data set. In each partition, we create training/test data as used in Goyal et al. (2020). Note however, different from Goyal et al. (2020), we scale the data in the range of $[-1, +1]$. In addition, here we generate uniform noise in range $[-1, +1]$ and use that as universum/contradiction samples. As seen from Table 1 the DOC³ outperforms all existing approaches (except adversarial based DROCC (Goyal et al., 2020) and PLAD for the Thyroid data); and significantly improves ($> 5-15\%$) upon the state-of-the-art results for the Arrhythmia and Abalone data. The optimal model parameters used for the results are provided in “Appendix B.1” (Table 7) for reproducibility. Note that, through out the paper we fix $\Delta = 0$. For all our experiments we see minimal improvements through tuning the Δ parameter. This is also discussed in our ablation studies in “Appendix C.1.2”.


5.2 Standard image benchmark datasets

5.2.1 CIFAR-10

For our next set of experiments we use the standard image benchmark CIFAR-10 dataset (Ruff et al., 2018; Goyal et al., 2020). The data consists of 32x32 colour images of 10 classes with 6000 images per class. The classes are mutually exclusive. The underlying task involves one-vs-rest anomaly detection, where we build a one class classifier for each class and evaluate it on the test data for all the 10-classes. Note that, this data does not have any naturally occurring universum (contradiction) samples (following Def. 2). So, we use synthetic universum samples by randomly generating the pixel values as $\sim \mathcal{N}(\mu, \sigma)$, with

Table 1 F1-score \pm standard deviation for one-vs-all anomaly detection on Thyroid, Arrhythmia, and Abalone datasets

Method	F1-Score Thyroid	Arrhythmia	Abalone
OC-SVM (Schölkopf et al., 1999)	0.39 \pm 0.01	0.46 \pm 0.00	0.48 \pm 0.00
DCN (Caron et al., 2018)	0.33 \pm 0.03	0.38 \pm 0.03	0.40 \pm 0.01
E2E-AE (Zong et al., 2018)	0.13 \pm 0.04	0.45 \pm 0.03	0.33 \pm 0.03
LOF (Breunig et al., 2000)	0.54 \pm 0.01	0.51 \pm 0.01	0.33 \pm 0.01
DAGMM (Zong et al., 2018)	0.49 \pm 0.04	0.49 \pm 0.03	0.20 \pm 0.03
DeepSVDD (Ruff et al., 2018)	0.73 \pm 0.00	0.54 \pm 0.01	0.62 \pm 0.01
GOAD (Bergman Hoshen, 2020)	0.72 \pm 0.01	0.51 \pm 0.02	0.61 \pm 0.02
DROCC (Goyal et al., 2020)	0.78 \pm 0.03	0.69 \pm 0.02	0.68 \pm 0.02
PLAD (Cai & Fan, 2022)	0.76 \pm 0.06	0.71 \pm 0.01	–
DOC ³ (ours)	0.74 \pm 0.01	0.73 \pm 0.01	0.77 \pm 0.01

Fig. 2 Random noise Universum (contradictions)

$\mu = 0$, $\sigma = 1$; where \mathcal{N} is the normal distribution (see Fig. 2). The idea of generating synthetic universum (contradiction) samples has been previously studied for binary (Weston et al., 2006; Cherkassky et al., 2011; Sinz et al., 2008), multiclass (Zhang & LeCun, 2017; Dhar et al., 2019) and regression (Dhar & Cherkassky, 2017) problems. In this paper we use such a similar mechanism for one class problems. Note that for the one-vs-rest AD problem, the generated universum samples do not belong to either ‘+1’ (normal) or ‘-1’ (anomalous) class used during testing (see Def. 2). The data is scaled in range $[-1, +1]$.

For this set of experiments we adopt a LeNet like architecture used in Ruff et al. (2018), Goyal et al. (2020). The detailed architecture specifics is provided in “Appendix B.1.1”. Note that, this paper focuses on the design and analysis of the DOC³ loss (Eq. 6). Here rather than adopting a state-of-the-art network architecture optimized for the specific dataset; we adopt a systematic approach to isolate the effectiveness of the proposed loss by using a basic LeNet architecture similar to Ruff et al. (2018), Goyal et al. (2020). This avoids secondary generalization effects encoded in most advanced architectures. To that end, the approaches in Ruff et al. (2018), Goyal et al. (2020) and DOC (Hinge in Eq. (3)) serve as the main baselines. In addition, for a more thorough comparison we also provide the results for DOC extended under disjoint auxiliary (DA) a.k.a. Outlier Exposure (OE) settings. For that we use

the additional universum samples as belonging to the negative class following (Goyal et al., 2020).

Table 2 provides the average \pm standard deviation of the AUC under the ROC curve over 10 runs of the experiment. Here, we report the results of the best performing DOC (Hinge in (3)) model selected over the range of parameters $\lambda = 1/2C = [1.0, 0.5]$ and that for DOC^3 over the range of parameters $\lambda = 1/2C = [0.1, 0.05]$, $C_U/C = [1.0, 0.5]$. We fix $\Delta = 0$. A more detailed discussion on model selection and the selected model parameters is provided in “Appendix B.3” for reproducibility. Note however, our results for the DROCC algorithm is different from that reported in Goyal et al. (2020). Re-running the codes provided in Goyal et al. (2020) did not yield similar results as reported in the paper (especially for ‘Ship’). Moreover, their current implementation normalizes the data using mean, $\mu = (0.4914, 0.4822, 0.4465)$ and standard deviation, $\sigma = (0.247, 0.243, 0.261)$. These values are calculated using the data from all the classes; which is not available during training of a single class. To avoid such inconsistencies we rather normalize using mean, $\mu = (0.5, 0.5, 0.5)$ and standard deviation, $\sigma = (0.5, 0.5, 0.5)$. Such a scale does not need apriori information of the other class’s pixel values and scales the data in a range of $[-1, +1]$. Detailed discussions on reproducing the results of the deep learning algorithms Deep-SVDD (Ruff et al., 2018) and DROCC (Goyal et al., 2020) is provided in “Appendix C.2” (see Tables 18, 19 and 20). As seen from Table 2, DOC^3 (using the noise universum), provides significant improvement ~ 5 –15% (and upto 30% for ‘Bird’), over its inductive counterpart (DOC). In addition, the DOC^3 in most cases outperforms the DOC (DA/OE). This illustrates the advantage of extending Anomaly Detection problems following Def. 2 in accordance with the Prop. 2.

Next, we show the effectiveness of extending the advanced adversarial based DROCC-LF method under universum settings over the OE based setting used in Goyal et al. (2020). The major difference is now the auxiliary data serves as universum samples and the loss function follows. (6) (see Algo. 1). For the DROCC-LF (OE) we use the same implementation as in Goyal et al. (2020). Additionally, we replace the relu operator $[x]_+$ with the softplus operator for the loss functions.

For our experiments, we adopt the same LeNet architecture used in Ruff et al. (2018), Goyal et al. (2020) (see “Appendix B.1.1”, Fig. 4). Finally, we run the experiments over 10 runs and report the best AUC over the range of parameters recommended in Goyal et al. (2020) (Sect. 5). That is learning rate = 10^{-4} , radius (r) in range of $\sqrt{d} = \{8.0, 16.0, 32.0\}$. Here, for both the methods we use Adam and fix the number of ascent steps = 10 and batch size = 256 and total epochs = 350. The remaining parameters are set to default values.

Table 2 provides the average \pm standard deviation of the AUC under the ROC curve over 10 runs of the experiment. We also provide the results for the standard DROCC (without any auxiliary data) (Goyal et al., 2020) and the more recent PLAD (an adversarial approach introduced to improve DROCC) (Cai & Fan, 2022) as baselines. As seen from Table 2 the DROCC-LF (univ) significantly outperforms the DROCC-LF (OE) method to upto - 30% (‘dog’) for some cases. Further, DROCC-LF (univ) outperforms the baseline algorithms for all cases except ‘Cat’. The final optimal parameters selected for the different classes is provided in “Appendix B.3”.

5.2.2 Fashion-MNIST (F-MNIST)

For our next set of experiments we use another standard image benchmark dataset F-MNIST (Xiao et al., 2017). The data consists of 28×28 gray images of Zalando’s fashion product

Table 2 Average AUC (with standard deviation) for one-vs-rest anomaly detection on CIFAR-10

CIFAR-10	Airplane	Automobile	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck
OC-SVM [†]	61.6±0.9	63.8±0.6	50.0±0.5	55.9±1.3	66.0±0.7	62.4±0.8	74.7±0.3	62.6±0.6	74.9±0.4	75.9±0.3
IF [†]	60.1 ± 0.7	50.8 ± 0.6	49.2±0.4	55.1±0.4	49.8±0.4	58.5±0.4	42.9±0.6	55.1±0.7	74.2±0.6	58.9±0.7
DCAE [†]	59.1 ± 5.1	57.4 ± 2.9	48.9±2.4	58.4±1.2	54.0±1.3	62.2±1.8	51.2±5.2	58.6±2.9	76.8±4.1	67.3±3.0
AnoGAN [†]	67.1 ± 2.5	54.7 ± 3.4	52.9±3.0	54.5±1.9	65.1±3.2	60.3±2.6	58.5±1.4	62.5±0.8	75.8±4.1	66.5±2.8
ConAD 16 [*]	77.2	63.1	63.1	61.5	63.3	58.8	69.1	64	75.5	63.7
1-NN [*]	69.02	44.2	68.27	51.32	76.71	49.97	72.44	51.13	69.09	43.33
Deep-SVDD (Soft) [†]	61.7 ± 4.2	64.8 ± 1.4	49.5±1.4	56.0±1.1	59.1±1.1	62.1±2.4	67.8±2.4	65.2±1.0	75.6±1.7	71.0±1.1
Deep-SVDD [†]	61.7 ± 4.1	65.9 ± 2.1	50.8±0.8	59.1±1.4	60.9±1.1	65.7±2.5	67.7±2.6	67.3±0.9	75.9±1.2	73.1±1.2
DOC (Hinge Eq. 3)	76.8 ± 1.4	62.6 ± 2.8	52.1±0.7	60.4±0.7	62.3±1.2	61.9±1.9	76.3±0.5	59.8±1.3	72.8±1.1	74.9±2.0
DOC (DA/OE)	69.5 ± 14.5	73.1 ± 3.3	67.3 ± 0.3	62.4 ± 2.5	71.1±2.1	67.3 ± 7.4	78.6 ± 1.7	66.8±2.4	70.3 ± 1.9	75.8 ± 2.5
DOC ³ (univ)	81.3 ± 0.5	74.2 ± 1.3	69.0 ± 0.6	62.1 ± 0.4	74.0±1.6	63.0±4.6	77.7 ± 0.3	67.6±1.8	81.1 ± 0.6	76.8 ± 2.0
PLAD ^{**}	82.5 ± 0.4	80.8 ± 0.9	68.8±1.2	65.2 ± 1.2	71.6±1.1	71.2±1.6	76.4±1.9	73.5±1.0	80.6±1.8	80.5±1.3
DROCC [‡]	79.2 ± 1.9	74.9 ± 2.6	68.3±1.5	62.3 ± 2.7	70.3±2.7	66.1±2.0	68.1±2.2	71.3±4.6	62.3±10.3	76.6±1.9
DROCC-LF(OE)	91.9 ± 0.9	70.5 ± 2.4	70.9 ± 2.2	63.1 ± 2.2	76.6±1.3	65.7±1.4	74.1 ± 2.7	70.6±3.6	85.1 ± 4.1	84.6 ± 2.4
DROCC-LF(univ)	96.8 ± 0.4	88.2 ± 5.3	79.8 ± 1.6	62.7 ± 5.3	80.4 ± 1.3	84.9 ± 4.1	87.4 ± 0.9	75.0 ± 1.4	93.4 ± 0.2	86.5 ± 0.8

[†]Results reported in Ruff et al. (2018). ^{*}Result reported in Goyal et al. (2020). ^{**}Results reported in Cai and Fan (2022). [‡]Our re-run of the algorithm in Goyal et al. (2020)

database and consists of 10 classes (product lines) with 60,000 training and 10,000 test samples. The classes are mutually exclusive. The underlying task involves one-vs-rest anomaly detection, where we build a one class classifier for each class and evaluate it on the test data for all the 10-classes. As before, this data does not have any naturally occurring universum (contradiction) samples (following Def. 2). So, we use synthetically generated universum samples from $\sim \mathcal{N}(\mu, \sigma)$, with $\mu = 0$, $\sigma = 1$; where \mathcal{N} is the normal distribution. The data is scaled in range $[-1, +1]$.

We adopt the same network and experiment set-up used in Cai and Fan (2022) (see Fig. 5 in “Appendix B.1.2”). As before, we provide the results for DOC and its extension under disjoint auxiliary (DA)/Outlier Exposure (OE) settings. We also provide the baseline results from Cai and Fan (2022). Table 3 provides the average \pm standard deviation of the AUC under the ROC curve over 10 runs of the experiment. A more detailed discussion on model selection and the selected model parameters is provided in “Appendix B.4” (see Tables 12, 13 and 14) for reproducibility.

As seen from Table 3, DOC³ (using the noise universum), provides significant improvement ~ 5 –20% over its inductive counterpart (DOC). In addition, the DOC³ in most cases outperforms the DOC (DA/OE). This further consolidates the advantage of extending Anomaly Detection problems under universum settings.

We also illustrate the effectiveness of extending the advanced adversarial based DROCC-LF method under universum settings over the OE based setting used in Goyal et al. (2020). We provide the results for the DROCC algorithm as baseline. In addition, we also provide the results of the recent PLAD (Cai & Fan, 2022) algorithm which adopts a generative adversarial learning approach proposed to stabilize DROCC. Note however, the results reported are from our re-run of the PLAD algorithm. We found several caveats with the code implementation and report these discrepancies in “Appendix C.3”.

As seen from Table 3 the DROCC-LF (univ) outperforms the DROCC-LF (OE) method and beats the baseline algorithms for all the classes. The final optimal parameters selected for the different classes is provided in “Appendix B.4” (see Table 15).

5.3 Visual inspection using real-life MV-Tec AD data

For our final set of experiments we tackle the more realistic visual inspection based anomaly detection problem in manufacturing lines. Lately with the recent advancements in deep learning technologies, there has been an increased interest towards automating manufacturing lines and adopting AI driven solutions providing automated visual inspection of product defects (Bergmann et al., 2019; Huang & Pan, 2015). One popular benchmark data set used for such problems is the MV-Tec AD data set (Bergmann et al., 2019).

5.3.1 Data set and experiment setup

The MV-Tec AD data set contains 5354 high-resolution color images of different industrial object and texture categories. For each categories it contains normal (no defect) images used for training. The test data contains both normal as well as anomalous (defective) product images. The anomalies manifest themselves in the form of over 70 different types of defects such as scratches, dents, contamination, and various other structural changes. The goal in this paper is to build one class image-level classifiers for the texture categories (see Table 4). We use the original data scale of $[0,1]$. Further, to simplify the problem we resize all the

Table 3 Average AUC (with standard deviation) for one-vs-rest anomaly detection on F-MNIST

F-MNIST	T-shirt	Trouser	Pullover	Dress	Coat	Sandal	Shirt	Sneaker	Bag	Ankle boot
OC-SVM [†]	86.1	93.9	85.6	85.9	84.6	81.3	78.6	97.6	79.5	97.8
IF [†]	91.0	97.8	87.2	93.2	90.5	93.0	80.2	98.2	88.7	95.4
Deep-SVDD [†]	79.1	94.0	83.0	82.9	87.0	80.3	74.9	94.2	79.1	93.2
DOC (Hinge Eq. 3)	71.01±7.2	87.2±4.4	77.0±3.5	83.2±2.7	78.6±2.1	95.8±1.2	62.6±3.1	91.1±3.8	71.5±2.3	91.3±2.9
DOC (DA/OE)	84.3±5.2	96.9±0.3	78.5±3.5	89.6±1.2	87.4±1.6	95.2±0.8	69.1±2.4	97.5±0.5	79.1±5.7	98.8±0.1
DOC ³ (univ)	87.0±2.4	97.3±0.4	84.8±2.6	91.4±0.7	89.2±1.5	96.0±2.3	74.6±1.3	97.7±0.6	78.9±1.9	98.2±0.4
PLAD [‡]	89.95±1.2	90.5±3.8	88.05±4.1	86.5±2.8	88.6±1.4	95.9±0.7	80.9±1.1	94.8±0.4	83.9±2.7	96.6±0.6
DROCC [†]	88.1	97.7	87.6	87.7	87.2	91.0	77.1	95.3	82.7	95.9
DROCC-LF(OE)	91.4±1.1	96.9±0.3	88.1±0.9	90.4±0.5	91.4±0.2	96.1±1.0	82.4±0.6	97.2±0.1	85.4±1.4	99.1±0.2
DROCC-LF(univ)	93.7±0.2	98.3±0.1	89.6±0.08	92.8±0.7	92.3±0.1	96.1±0.7	82.1±0.2	98.0±0.1	87.9±0.4	99.3±0.05

[†]Results reported in Cai and Fan (2022). [‡]Our re-run of the algorithm in Cai and Fan (2022)

Table 4 MVTec-AD dataset

Textures	Train	Test	
		Normal	Anomaly
Carpet	280	28	89
Leather	245	32	92
Tile	230	33	84
Wood	247	19	60

images to 64×64 pixel. Note that, for the current analysis we only use the texture classes containing RGB images.

For this problem we have naturally occurring universum (contradiction) samples in the form of the objects' images or other texture types. That is, for the goal of building a one class classifier for 'carpet', all the 'other textures' (leather, tile, wood) or the 'objects' (bottle, cable, capsule, hazelnut, metal nut, pill, transistor) available in the dataset, can serve as universum (contradiction) samples. This is inline with the problem setting in Def. 2, where such samples are neither 'normal' nor 'anomalous' (defective) carpet samples. For our experiments, we use three types of universum,

- *Noise*: Similar to previous experiments we generate random noise as universum samples. Here, since the data is already scaled in the range of $[0, 1]$, we generate 64×64 dimension images where the pixel values are obtained from a uniform distribution $\sim \mathcal{U}(0, 1)$.
- *Objects*: This type of universum contains all the images in the object categories with RGB pixels viz. bottle, cable, capsule, hazelnut, metal nut, pill, transistor. Note that, we include both the normal as well as the defective samples for these objects.
- *Other Textures*: Here we use the remaining texture images as universum. That is, if the goal is building a one class classifier for 'carpet' we use the images from the other 'textures' (leather, tile, wood) as universum. We include both the normal as well as the defective samples in the universum set.

As before, we adopt a LeNet like architecture (schematic representation in Fig. 3, details in "Appendix B.1.3", Fig. 6). Note that, there have been a few recent works proposing advanced architectures to achieve state-of-the-art performance on this data (Carrara et al., 2020; Huang et al., 2019). However, the main focus here is to isolate the effectiveness of DOC³, and hence we mainly compare against DOC and DOC(OE) baselines using a simple LeNet network. Since our baselines DOC, DOC(OE) using LeNet have not been previously reported on this data; as sanity check we also add the results in Massoli et al. (2020) for a good comparison with different classes of algorithms. Also, we adopt a slight modification to our loss function. Rather than using relu function $[x]_+$ in Eq. (3), and (6) for the training samples; we use a softplus operator. We see improved results using this modification. Note that, softplus is a dominating surrogate loss over relu, and hence Theorem 1 still holds.

5.3.2 Performance comparison results

Table 5 provides the results over 10 runs of our experiments. We provide the the average \pm standard deviation of the AUC values for DOC, DOC (DA/OE) and DOC³ algorithm. In addition we also provide the best AUC obtained for each algorithm over these 10 runs. Additional details on model selection and the optimal hyperparameters is provided in "Appendix B.5". As seen in Table 5, the DOC³ algorithm provides significant improvement over DOC.

Table 5 AUC for MV/Tec-AD (Texture) data

TEXTURES	AE [‡] _{L2}	Geo Trans [‡]	GANomaly [‡]	ITAE [‡]	EGBAD [‡]	CBI-GAN [‡]	DOC		DOC (OE)		DOC (noise)		DOC ³ (noise)		DOC ³ (objects)		DOC ³ (textures)		
							Eq. (3) Best	Avg. ± std	Best	Avg. ± std	Best	Avg. ± std	Best	Avg. ± std	Best	Avg. ± std	Best	Avg. ± std	Best
Carpet	64	44	70	71	52	55	81.1	76.2	89.6	54.9	95.7	93.8	81.1	81.1	81.1	81.1	81.1	81.1	81.1
							81.1 ± 0.0	56.5 ± 10.1	82.1 ± 4.2	49.2 ± 4.9	80.4 ± 8.4	87.5 ± 3.8	81.1 ± 0.0	81.1 ± 0.0	81.1 ± 0.0	81.1 ± 0.0	81.1 ± 0.0	81.1 ± 0.0	81.1 ± 0.0
Leather	80	84	84	86	55	83	63.1	65.7	95.5	40.1	88.1	93.5	63.1	63.1	63.1	63.1	63.1	63.1	63.1
							62.7 ± 0.3	64.8 ± 0.8	89.6 ± 5.1	39.8 ± 0.2	82.9 ± 4.5	83.1 ± 7.5	62.4 ± 0.5	62.4 ± 0.5	62.4 ± 0.5	62.4 ± 0.5	62.4 ± 0.5	62.4 ± 0.5	62.4 ± 0.5
Tile	74	42	79	74	79	91	62.8	65.9	75.7	67.7	66.3	77.0	65.1	65.1	65.1	65.1	65.1	65.1	65.1
							62.3 ± 0.7	64.9 ± 0.7	74.0 ± 1.4	65.3 ± 0.9	64.7 ± 0.6	76.5 ± 0.5	64.4 ± 0.5	64.4 ± 0.5	64.4 ± 0.5	64.4 ± 0.5	64.4 ± 0.5	64.4 ± 0.5	64.4 ± 0.5
Wood	97	61	83	92	91	95	41.1	90.2	77.2	52.6	93.1	75.3	49.4	49.4	49.4	49.4	49.4	49.4	49.4
							40.6 ± 0.1	82.8 ± 5.6	70.9 ± 6.7	50.5 ± 1.5	83.4 ± 7.0	69.0 ± 5.8	49 ± 0.4	49 ± 0.4	49 ± 0.4	49 ± 0.4	49 ± 0.4	49 ± 0.4	49 ± 0.4

[‡] Results taken from Massoli et al. (2020). Bold = best overall model. Underline = best universum or OE model




Fig. 3 Schematic representation of the Network used for MVTec-AD results in Table 5

Depending on the type of universum typical improvements range upto $> 50\%$. In addition, DOC^3 provides consistent improvements over the DOC (DA/OE) algorithm. In all, these results further consolidate the utility of DOC^3 under the universum setting (Def 2). Separately, Table 5 also provides the baseline results available in Massoli et al. (2020). Note that, these results are obtained using advanced network architectures adopted for the MVTec data, and are not averaged over multiple runs. Hence, we compare these results with the best AUC obtained for DOC, DOC (DA/OE) and DOC^3 over 10 runs. As seen from Table 5, DOC^3 improves upon the ‘carpet’ and ‘leather’ results using the ‘objects’ universum. Further, it achieves comparable performance for the ‘Wood’ and ‘Tile’ texture using ‘Noise’ and ‘Obj.’ universum respectively. Achieving improved performance over the baseline algorithms, even using a basic LeNet architecture sheds a very positive note for the proposed DOC^3 algorithm.

5.3.3 Understanding DOC^3 performance using Theorem (2)

For our final set of experiments we try to understand the working of the DOC^3 algorithm in connection with the correlation $\Sigma(\infty)$ (in Theorem 2). Table 6 reports the correlation values for the training and universum samples using ‘RAW’ pixel, ‘DOC’ and DOC^3 solution’s feature maps. For the feature map we use the CNN features shown in Fig. 3. Also, the DOC^3 solutions represent the estimated model using the training data (in column 1) and the respective universum data (in column 2). As seen from the results, the DOC solution provides high correlation $\Sigma(\infty)$ between the training and universum samples. In essence, the DOC solution sees the training and universum samples similarly. This is not desirable, as the universum samples follow a different distribution than training samples. On the contrary, the DOC^3 provides a solution where the correlation between the training and universum samples are significantly reduced. This is inline with the Theorem 2’s analysis (Sect. 3.2), where we argued that the DOC^3 searches for a solution with low $\Sigma(\infty)$ between the training and universum samples (in feature space). And by doing so ensures lower ERC and improved generalization compared to DOC (confirmed empirically in Table 5). Another interesting point seen for the ‘other texture’ universum type, with originally high raw pixel correlation values (~ 0.9) is that; using DOC^3 provides limited improvement. Such universum types are too similar to the training data, and act as ‘bad’ contradictions.

6 Future research

Broadly there are two major future research directions,

Table 6 Average \pm standard deviation of correlation $\Sigma(\infty)$ between training and universum over 10 runs

Train Data	Univ. Data	Raw	DOC Feat. map	DOC ³ Feat. map
Carpet	Noise	73.8	99.9 \pm 0.0	17.5 \pm 3.8
	Obj	73.8	97.1 \pm 0.1	28.8 \pm 3.4
	Text	92.4	99.9 \pm 0.0	68.7 \pm 0.0
Leather	Noise	70.1	99.3 \pm 0.1	42.1 \pm 1.9
	Obj	70.7	91.9 \pm 0.3	23.5 \pm 4.3
	Text	93.6	99.3 \pm 0.1	97.0 \pm 1.7
Tile	Noise	71.2	99.8 \pm 0.0	54.2 \pm 1.0
	Obj	71.5	92.8 \pm 0.8	31.3 \pm 8.8
	Text	89.9	99.6 \pm 0.0	89.3 \pm 2.2
Wood	Noise	69.5	99.8 \pm 0.0	39.8 \pm 5.8
	Obj	69.8	93.5 \pm 0.2	41.9 \pm 7.1
	Text	91.7	99.7 \pm 0.0	64.4 \pm 2.1

Values scaled $\times 10^2$

Model selection This is a generic issue for any (unsupervised) one class based anomaly detection formulation, and is further complicated by the non-convex loss landscape for deep learning problems. For DOC³ we simplify model selection by fixing $\Delta = 0$, and optimally tuning C_U . However, the success of DOC³ heavily depends on carefully tuning of its hyperparameters. In the absence of any validation set containing both ‘normal’ and ‘anomalous’ samples, we follow the current norm of reporting the best model’s results over a small subset of hyperparameters. But this is far from practical. We believe, our Theorem 1 provides a good framework for bound based model selection. This in conjunction with Theorem 2 and the recent works on ERC for deep architectures (Neishabur et al., 2015; Sokolic et al., 2016), may provide better mechanisms for model selection and yield optimal models.

Selecting ‘good’ universum samples The effectiveness of DOC³ also depends on the type of universum used. Our analysis in Sect. 5.3.3 provides some initial insights into the workings of DOC³, and how to loosely identify ‘bad’ contradictions. Additional analysis, possibly inline with the Histogram of Projections (HOP) technique introduced in Cherkassky et al. (2011), Dhar et al. (2019), is needed to improve our understanding of ‘good’ universum samples. This is an open research problem.

7 Conclusions

This paper introduces the notion of learning from contradictions for deep one class classification and introduces the DOC³ algorithm. DOC³ is shown to provide improved generalization over DOC, its inductive counterpart, by deriving the Empirical Rademacher Complexity (ERC). We empirically show the effectiveness of the proposed formulation, and connect the results to our theoretical analysis. Finally, we also discuss the limitations and the future research directions.

Author contributions A detailed list of contribution is provided below, Conceptualization—SD Data curation—BGT Formal Analysis—SD Funding acquisition—NA Investigation—SD, BGT Methodology—SD, BGT Project administration—SD Resources—SD, BGT Software—BGT, SD Supervision—SD. Validation—BGT Visualization—SD, BGT. Writing (original draft)—SD Writing (review, editing)—SD, BGT.

Funding Not applicable.

Data availability All data used is publicly available and appropriately referenced in the paper.

Declarations

Conflict of interest Not applicable.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Code availability All our codes available at <https://github.com/sauptikdhar/DOC3>

Appendix A Proofs

A.1 Proof of Proposition 1

Part 1 A slightly different version of this proposition is analyzed in Proposition 8.2 of Scholkopf et al. (2002) and Chang and Lin (2001). Here, we provide a different version of the connection between the solutions of (3) and (2). This is achieved through analyzing the KKT systems of the formulations. We start with the formulation (3). Note that, (3) is the same as solving,

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^n \xi_i \\ \text{s.t.} \quad & \mathbf{w}^\top \phi(\mathbf{x}_i) \geq 1 - \xi_i; \quad \xi_i \geq 0 \end{aligned} \quad (\text{A1})$$

The Lagrangian is given as,

$$\mathcal{L}(\mathbf{w}, \xi, \alpha, \beta) = \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^n \xi_i - \sum_{i=1}^n \beta_i \xi_i - \sum_{i=1}^n \alpha_i [\mathbf{w}^\top \phi(\mathbf{x}_i) - 1 + \xi_i]$$

KKT System

$$\nabla_{\mathbf{w}} \mathcal{L} = 0 \Rightarrow \mathbf{w} = \sum_{i=1}^n \alpha_i \phi(\mathbf{x}_i) \quad (\text{A2})$$

$$\nabla_{\xi} \mathcal{L} = 0 \Rightarrow C = \alpha_i + \beta_i \quad (\text{A3})$$

Complimentary Slackness,

$$\alpha_i [\mathbf{w}^\top \phi(\mathbf{x}_i) - 1 + \xi_i] = 0 \quad (\text{A4})$$

$$\beta_i \xi_i = 0 \quad (\text{A5})$$

Constraints,

$$\mathbf{w}^\top \phi(\mathbf{x}_i) \geq 1 - \xi_i \quad (\text{A6})$$

$$\xi_i \geq 0 \quad (\text{A7})$$

Define $\delta = \frac{1}{\sum_i \alpha_i}$ and re-write the equations (A2)–(A7) by scaling with $\delta > 0$ as, $\hat{\mathbf{w}} = \mathbf{w}\delta$; $\hat{\alpha}_i = \alpha_i\delta$; $\hat{\beta}_i = \beta_i\delta$; $\hat{\xi}_i = \xi_i\delta$. ($\delta > 0$; $\because \exists i$ s.t. $\alpha_i > 0$ and $\forall i, \alpha_i \geq 0$)

This gives, Transformed KKT System

$$\hat{\mathbf{w}} = \sum_{i=1}^n \hat{\alpha}_i \phi(\mathbf{x}_i) \quad (\text{A8})$$

$$C\delta = \hat{\alpha}_i + \hat{\beta}_i \quad (\text{A9})$$

Complimentary Slackness,

$$\hat{\alpha}_i [\hat{\mathbf{w}}^\top \phi(\mathbf{x}_i) - \delta + \hat{\xi}_i] = 0 \quad (\text{A10})$$

$$\hat{\beta}_i \hat{\xi}_i = 0 \quad (\text{A11})$$

Constraints,

$$\hat{\mathbf{w}}^\top \phi(\mathbf{x}_i) \geq \delta - \hat{\xi}_i; \quad \hat{\xi}_i \geq 0 \quad (\text{A12})$$

Note that, the transformed KKT system (A8)–(A12) solves (2) with $\nu = \frac{1}{Cn\delta}$; $\rho = \delta$ (compare with the KKT of (2)).

Part 2 For the solution to (2) obtained from Proposition 1 (i) the decision rule can be given as,

$$\begin{aligned} \hat{\mathbf{w}}^\top \phi(\mathbf{x}) - \rho = 0 &\Rightarrow (\mathbf{w}\delta)^\top \phi(\mathbf{x}) - \delta = 0 \\ &\Rightarrow \mathbf{w}^\top \phi(\mathbf{x}) - 1 = 0 \quad (\because \delta > 0) \end{aligned}$$

□

A.2 Proof of Proposition 2

Note that for this proof we need to accommodate a case where a sample may not belong to either of the two classes $\{-1, +1\}$. For this we rather analyze a different decision rule than (3).

Define,

$$g(\mathbf{x}) = \begin{cases} +1; & \text{if } \mathbf{w}^\top \phi(\mathbf{x}_i) > 1 \\ -1; & \text{if } \mathbf{w}^\top \phi(\mathbf{x}_i) < 1 \end{cases}$$

This gives,

$$\begin{aligned} g(\mathbf{x}) = +1 &\Rightarrow h(\mathbf{x}) = +1 \text{ and } g(\mathbf{x}) = -1 \Rightarrow h(\mathbf{x}) = -1 \\ &\Rightarrow \mathbf{P}_{\mathcal{D}_{\mathcal{U}}}(h(\mathbf{x}^*) = y) \geq \mathbf{P}_{\mathcal{D}_{\mathcal{U}}}(g(\mathbf{x}^*) = y) \quad \forall y = \{-1, +1\} \end{aligned}$$

Since the events are mutually exclusive we have,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[h(\mathbf{x}^*) = y] \right) &\geq \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[g(\mathbf{x}^*) = y] \right) \\ &\Rightarrow \inf_h \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[h(\mathbf{x}^*) = y] \right) \geq \inf_g \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[g(\mathbf{x}^*) = y] \right) \end{aligned}$$

$$\begin{aligned} \Rightarrow 1 - \inf_h \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[h(\mathbf{x}^*) = y] \right) &\leq 1 - \inf_g \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcup_{y \in \{-1, +1\}} \mathbf{1}[g(\mathbf{x}^*) = y] \right) \\ \Rightarrow \sup_h \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcap_{y \in \{-1, +1\}} \mathbf{1}[h(\mathbf{x}^*) \neq y] \right) &\text{ (De-Morgan's law)} \\ &\leq \sup_g \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} \left(\bigcap_{y \in \{-1, +1\}} \mathbf{1}[g(\mathbf{x}^*) \neq y] \right) \\ &= \sup_g \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} (\mathbf{1}[g(\mathbf{x}^*) \neq +1] \wedge \mathbf{1}[g(\mathbf{x}^*) \neq -1]) \\ &= \sup_g \mathbb{E}_{\mathcal{D}_{\mathcal{U}}} (\mathbf{1}[\mathbf{w}^\top \phi(\mathbf{x}) - 1 \not\geq 0] \wedge \mathbf{1}[\mathbf{w}^\top \phi(\mathbf{x}) - 1 \not\leq 0]) \end{aligned}$$

The maximum can be achieved when $\mathbf{w}^\top \phi(\mathbf{x}) - 1 = 0$ □

A.3 Proof of Theorem 1

Define, $R_{f, \theta} = \{\mathbf{x} : f(\mathbf{x}) \geq \theta\}$. This gives,

$$\mathbf{P}_{\mathcal{D}_{\mathcal{T}}|\mathcal{Y}=\pm 1} \{\mathbf{x} \notin R_{f, 1-\kappa}\} = \mathbb{E}_{\mathcal{D}_{\mathcal{T}}|\mathcal{Y}=\pm 1} [H(f(\mathbf{x}), 1 - \kappa)] \tag{A13}$$

where, $H(x, \theta) = \begin{cases} 0; & \text{if } x \geq \theta \\ 1; & \text{else} \end{cases}$. For the rest of the proof we drop the subscripts as it is clear from context. To bound the R.H.S of (A13) we follow a similar approach of bounding a dominating function see Theorem 4.17 in Shawe-Taylor et al. (2004). Here we define,

$$A(x) = \begin{cases} 0; & \text{if } x > 1 \\ \frac{1-x}{\kappa}; & \text{if } 1 - \kappa \leq x \leq 1 \\ 1; & \text{if } x < 1 - \kappa \end{cases}$$

Note that, $A(x)$ is $\frac{1}{\kappa}$ -Lipchitz. Further, $H(f(\mathbf{x}), 1 - \kappa) \leq A(f(\mathbf{x}))$. This gives, $\mathbb{E}[H(f(\mathbf{x}), 1 - \kappa) - 1] \leq \mathbb{E}[A(f(\mathbf{x})) - 1]$. Hence with probability $1 - \eta, \forall f \in \mathcal{F}$ the following holds (see Theorem 4.9 in Shawe-Taylor et al. (2004)); where $\hat{\mathbb{E}}$ = the empirical estimate for the expectation operator.

$$\mathbb{E}[H(f(\mathbf{x}), 1 - \kappa) - 1] \leq \hat{\mathbb{E}}[A(f(\mathbf{x})) - 1] + \hat{\mathcal{R}}_n((A - 1) \circ \mathcal{F}) + 3\sqrt{\frac{\ln \frac{2}{\eta}}{2n}}$$

From Th. 4.15 (Shawe-Taylor et al., 2004)

$$\Rightarrow \mathbb{E}[H(f(\mathbf{x}), 1 - \kappa)] \leq \frac{\sum_{i=1}^n \xi_i}{\kappa n} + \frac{2}{\kappa} \hat{\mathcal{R}}_n(\mathcal{F}) + 3\sqrt{\frac{\ln \frac{2}{\eta}}{2n}}$$

where, $\xi_i = [1 - f(\mathbf{x}_i)]_+$. Using (A13), we get the final form of Theorem (1). □

A.4 Proof of Theorem 2

Part 1: It is clear that $\mathcal{F}_{\text{univ}} \subseteq \mathcal{F}_{\text{ind}}$. This ensures $\hat{\mathcal{R}}_n(\mathcal{F}_{\text{univ}}) \leq \hat{\mathcal{R}}_n(\mathcal{F}_{\text{ind}})$ (following Theorem 4.15 (i) in Shawe-Taylor et al. (2004)). \square

Part 2(a): This follows from standard analysis (see Theorem 4.12 (Shawe-Taylor et al., 2004) or Lemma 22 in Bartlett and Mendelson (2002)).

$$\begin{aligned}
 \hat{\mathcal{R}}_n(\mathcal{F}_{\text{ind}}) &= \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right| \middle| (\mathbf{x}_i)_{i=1}^n \right] \\
 &= \mathbb{E}_\sigma \left[\sup_{\|\mathbf{w}\|^2 \leq \Lambda^2} \left| \frac{2}{n} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right| \middle| (\mathbf{x}_i)_{i=1}^n \right] \\
 &= \mathbb{E}_\sigma \left[\sup_{\|\mathbf{w}\|^2 \leq \Lambda^2} \left| \frac{2}{n} \mathbf{w}^\top \left(\sum_{i=1}^n \sigma_i \phi(\mathbf{x}_i) \right) \right| \middle| (\mathbf{x}_i)_{i=1}^n \right] \\
 &\leq \mathbb{E}_\sigma \left[\sup_{\|\mathbf{w}\|^2 \leq \Lambda^2} \frac{2\|\mathbf{w}\|}{n} \left\| \sum_{i=1}^n \sigma_i \phi(\mathbf{x}_i) \right\| \middle| (\mathbf{x}_i)_{i=1}^n \right] \\
 &\leq \frac{2\Lambda}{n} \mathbb{E}_\sigma \left[\left\| \sum_{i=1}^n \sigma_i \phi(\mathbf{x}_i) \right\| \middle| (\mathbf{x}_i)_{i=1}^n \right] \\
 &\leq \frac{2\Lambda}{n} \mathbb{E}_\sigma \left[\left\| \sum_{i=1}^n \sigma_i \phi(\mathbf{x}_i) \right\|^2 \middle| (\mathbf{x}_i)_{i=1}^n \right]^{\frac{1}{2}} \quad (\text{Jensen's inequality}) \\
 &\leq \frac{2\Lambda}{n} \mathbb{E}_\sigma \left[\sum_{i,j=1}^n \sigma_i \sigma_j \phi(\mathbf{x}_i) \phi(\mathbf{x}_j) \middle| (\mathbf{x}_i)_{i=1}^n \right]^{\frac{1}{2}} \\
 &= \frac{2\Lambda}{n} \left[\sum_{i=1}^n \|\phi(\mathbf{x}_i)\|^2 \right]^{\frac{1}{2}}
 \end{aligned}$$

\square

Part 2(b): Define

$$\begin{aligned}
 \mathcal{W}_{\text{univ}} &= \{ \mathbf{w} \mid \|\mathbf{w}\|^2 \leq \Lambda^2; |\mathbf{w}^\top \phi(\mathbf{x}^*) - 1| \leq \Delta; \forall \mathbf{x}^* \in \mathcal{X}_U^* \} \\
 &\subseteq \{ \mathbf{w} \mid \|\mathbf{w}\|^2 \leq \Lambda^2; |\mathbf{w}^\top \mathbf{u}_j - 1| \leq \Delta; \\
 &\quad \forall \mathbf{u}_j = \phi(\mathbf{x}_j^*); \forall \mathbf{x}_j^* \in \mathcal{X}_U^*; j = 1 \dots m \}
 \end{aligned} \tag{A14}$$

\therefore the constraint on all $\mathbf{x}^* \in \mathcal{X}_U^* \Rightarrow$ constraint on $m-$ samples. Now, let's analyze the constraint $|\mathbf{w}^\top \mathbf{u}_j - 1| \leq \Delta$. This implies, $\mathbf{w}^\top \mathbf{u}_j - 1 \leq \Delta; 1 - \mathbf{w}^\top \mathbf{u}_j \leq \Delta$ (simultaneously). However, only one of the constraint is active. Hence, we re-write the constraint as,

$$\forall j; \begin{bmatrix} \mathbf{w}^\top \mathbf{u}_j \\ \mathbf{w}^\top (-\mathbf{u}_j) \end{bmatrix} \leq \begin{bmatrix} \Delta + 1 \\ \Delta - 1 \end{bmatrix}.$$

Next define a mapping where we concatenate the reflected space. i.e. $\psi : \phi(\mathbf{x}^*) \rightarrow$

$$\begin{bmatrix} \phi(\mathbf{x}^*)^\top \\ -\phi(\mathbf{x}^*)^\top \end{bmatrix} \text{ and rewrite } V = \psi\left([\phi(\mathbf{x}_j^*)]_{j=1}^m\right) = \begin{bmatrix} \phi(\mathbf{x}_1^*)^\top \\ \phi(\mathbf{x}_2^*)^\top \\ \vdots \\ \phi(\mathbf{x}_m^*)^\top \\ -\phi(\mathbf{x}_1^*)^\top \\ -\phi(\mathbf{x}_2^*)^\top \\ \vdots \\ -\phi(\mathbf{x}_m^*)^\top \end{bmatrix}. \text{ This can be compactly}$$

re-written as, $V = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \otimes \begin{bmatrix} (\mathbf{u}_1)^T \\ \vdots \\ (\mathbf{u}_m)^T \end{bmatrix}$. This results to the overall constraint in (A14) to be,

$$\mathcal{W}_{\text{univ}} \subseteq \{\mathbf{w} \mid \|\mathbf{w}\|^2 \leq \Delta^2; [V\mathbf{w}]_j \leq \epsilon_j; j = 1 \dots 2m\}$$

where, $\epsilon = \begin{bmatrix} \Delta + 1 \\ \vdots \\ \Delta + 1 \\ \Delta - 1 \\ \vdots \\ \Delta - 1 \end{bmatrix} \begin{matrix} \left. \vphantom{\begin{matrix} \Delta + 1 \\ \vdots \\ \Delta + 1 \\ \Delta - 1 \\ \vdots \\ \Delta - 1 \end{matrix}} \right\} \text{m-times} \\ \left. \vphantom{\begin{matrix} \Delta + 1 \\ \vdots \\ \Delta + 1 \\ \Delta - 1 \\ \vdots \\ \Delta - 1 \end{matrix}} \right\} \text{m-times} \end{matrix}$.

In essence for each constraint in (A14) we create $2 \times$ the constraints for both the original and reflected space to take care of the absolute value.

Now,

$$\begin{aligned} \mathcal{W}_{\text{univ}} &\subseteq \{\mathbf{w} \mid \|\mathbf{w}\|^2 \leq \Delta^2; [V\mathbf{w}]_j \leq \epsilon_j \forall j = 1 \dots 2m\} \\ &\subseteq \{\mathbf{w} \mid \|\mathbf{w}\|^2 \leq \Delta^2; (\mathbf{w}^\top V^\top V \mathbf{w}) \leq 2m[\Delta^2 + 1]\} \end{aligned} \tag{A15}$$

The last line follows as the element-wise constraint is relaxed by $\|\cdot\|_2^2$ constraint.

Next, from (A15) and assuming a fixed mapping $\phi(\cdot)$, for the given training data $Z =$

$$\begin{bmatrix} (\mathbf{z}_1)^\top \\ \vdots \\ (\mathbf{z}_n)^\top \end{bmatrix} = \begin{bmatrix} \phi(\mathbf{x}_1)^\top \\ \vdots \\ \phi(\mathbf{x}_n)^\top \end{bmatrix} \text{ we have,}$$

$$\begin{aligned} &\hat{\mathcal{R}}(\mathcal{F}_{\text{univ}}) \\ &\stackrel{(a)}{=} \frac{2}{n} \mathbb{E}_\sigma \left[\sup_{\mathbf{w} \in \mathcal{W}_{\text{univ}}} \sigma^\top(Z\mathbf{w}) \right] \\ &\leq \frac{2}{n} \mathbb{E}_\sigma \left[\sup_{\substack{\|\mathbf{w}\|^2 \leq \Delta^2 \\ (\mathbf{w}^\top V^\top V \mathbf{w}) \leq 2m[\Delta^2 + 1]}} \sigma^\top(Z\mathbf{w}) \right] \text{ (from (A15))} \end{aligned}$$

Hence $\forall \gamma \geq 0$ and $\Gamma = \Delta^2 + 2\gamma m(\Delta^2 + 1)$ we have,

$$\stackrel{(b)}{\leq} \frac{2}{n} \mathbb{E}_\sigma \left[\sup_{\|\mathbf{w}\|^2 + \gamma(\mathbf{w}^\top V^\top V \mathbf{w}) \leq \Gamma} \sigma^\top(Z\mathbf{w}) \right]$$

$$\begin{aligned}
 &= \frac{2}{n} \mathbb{E}_\sigma \left[\sup_{\mathbf{w}^\top \left(\frac{I + \gamma V^\top V}{\Gamma} \right) \mathbf{w} \leq 1} \boldsymbol{\sigma}^\top (Z\mathbf{w}) \right] \\
 &\stackrel{(c)}{=} \frac{2}{n} \mathbb{E}_\sigma \left[\left\| \left(\frac{I + \gamma V^\top V}{\Gamma} \right)^{-\frac{1}{2}} Z^\top \boldsymbol{\sigma} \right\| \right] \\
 &\leq \frac{2}{n} \left[\mathbb{E}_\sigma \left[\left\| \left(\frac{I + \gamma V^\top V}{\Gamma} \right)^{-\frac{1}{2}} Z^\top \boldsymbol{\sigma} \right\|^2 \right] \right]^{\frac{1}{2}} \text{ (Jensen's inequality)} \tag{A16}
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(d)}{=} \frac{2}{n} \left[\text{tr} \left[Z \left(\frac{I + \gamma V^\top V}{\Gamma} \right)^{-1} Z^\top \right] \right]^{\frac{1}{2}} \text{ (tr := Trace)} \\
 &= \frac{2}{n} \left[\Lambda^2 + 2\gamma m(\Delta^2 + 1) \right]^{\frac{1}{2}} \left[\text{tr} \left[Z(I + \gamma V^\top V)^{-1} Z^\top \right] \right]^{\frac{1}{2}} \\
 &\stackrel{(e)}{=} \left(\frac{2}{n} \right) (\Lambda^2 + 2\gamma m(\Delta^2 + 1))^{\frac{1}{2}} \left[\text{tr}(Z^\top Z) \right. \\
 &\quad \left. - \gamma \text{tr} \left[ZV^\top (I + \gamma VV^\top)^{-1} VZ^\top \right] \right]^{\frac{1}{2}} \text{ (}\forall \gamma \geq 0\text{)} \\
 &= \frac{2\Lambda}{n} \sqrt{\left\| \sum_{i=1}^n \phi(\mathbf{x}_i) \right\|^2 \left[\left(1 + \frac{2\gamma m(\Delta^2 + 1)}{\Lambda^2} \right) \right.} \\
 &\quad \left. \left[1 - \frac{\gamma \text{tr} \left[(I + \gamma VV^\top)^{-1} (VZ^\top ZV^\top) \right]}{\text{tr}(Z^\top Z)} \right] \right]^{\frac{1}{2}} \\
 &\Rightarrow \hat{\mathcal{R}}(\mathcal{F}_{\text{univ}}) \stackrel{(f)}{\leq} \frac{2\Lambda}{n} \sqrt{\left\| \sum_{i=1}^n \phi(\mathbf{x}_i) \right\|^2 \left[\left(1 + \frac{2\gamma m(\Delta^2 + 1)}{\Lambda^2} \right) \right.} \\
 &\quad \left. \left[1 - \frac{\gamma \text{tr}(VZ^\top ZV^\top)}{\text{tr}(I + \gamma VV^\top) \text{tr}(Z^\top Z)} \right] \right]^{\frac{1}{2}}
 \end{aligned}$$

The (in)-equalities follow,

- a. from symmetry $\mathbf{w} \in \mathcal{W}_{USVM} \Rightarrow -\mathbf{w} \in \mathcal{W}_{USVM}$. Hence we drop the absolute term from definition. Also for simplicity we drop the conditional term. This is clear from context.
- b. since the conditions $\frac{\|\mathbf{w}\|^2 \leq \Lambda^2}{(\mathbf{w}^\top V^\top V \mathbf{w}) \leq 2m[\Delta^2 + 1]} \Rightarrow \|\mathbf{w}\|^2 + \gamma(\mathbf{w}^\top V^\top V \mathbf{w}) \leq \Gamma \quad \forall \gamma \geq 0$.
- c. stationary point of the constraint. A similar approach was previously used in Rosenberg and Bartlett (2007).
- d. since (Rademacher variables) are drawn uniformly over $\sigma \sim \{-1, +1\}$; we cancel the cross-terms $\sigma_i \sigma_j$ under expectation \mathbb{E}_σ .
- e. using Sherman-Morrison-Woodbury formula.
- f. from the matrix inequality II in Patel and Toda (1979).

□

Layer (type)	Output Shape	Param #
Conv2d-1	[-1, 32, 32, 32]	2,400
BatchNorm2d-2	[-1, 32, 32, 32]	0
MaxPool2d-3	[-1, 32, 16, 16]	0
Conv2d-4	[-1, 64, 16, 16]	51,200
BatchNorm2d-5	[-1, 64, 16, 16]	0
MaxPool2d-6	[-1, 64, 8, 8]	0
Conv2d-7	[-1, 128, 8, 8]	204,800
BatchNorm2d-8	[-1, 128, 8, 8]	0
MaxPool2d-9	[-1, 128, 4, 4]	0
Linear-10	[-1, 128]	262,144
Linear-11	[-1, 64]	8,192
Linear-12	[-1, 1]	64

Fig. 4 Network used for CIFAR10 results in Table 2

Layer (type)	Output Shape	Param #
Conv2d-1	[-1, 16, 28, 28]	400
BatchNorm2d-2	[-1, 16, 28, 28]	0
MaxPool2d-3	[-1, 16, 14, 14]	0
Conv2d-4	[-1, 32, 14, 14]	12,800
BatchNorm2d-5	[-1, 32, 14, 14]	0
MaxPool2d-6	[-1, 32, 7, 7]	0
Linear-7	[-1, 128]	200,704
Linear-8	[-1, 64]	8,192
Linear-9	[-1, 1]	64

Fig. 5 Network used for FMNIST results in Table 2

Appendix B Reproducibility

B.1 Network architectures

B.1.1 LeNet architecture for CIFAR-10 experiments

For CIFAR-10 we use the same architecture (Fig. 4) as used in Goyal et al. (2020).

B.1.2 LeNet architecture for FMNIST experiments

For FMNIST we use the same architecture as described in Cai and Fan (2022), see Fig. 5. Note however, there are discrepancies between the official code <https://openreview.net/forum?id=-Xdts90bWZ3> and the description in the paper. We adopt the description in paper; as the implementation is inconsistent with general deep learning theory. For additional details on these issues please see “Appendix C.3”.

B.1.3 LeNet architecture for MVTEC experiments

For MVTEC-AD there have been a few recent works proposing advanced architectures to achieve state-of-the-art performance on this data (Carrara et al., 2020; Huang et al., 2019). However, the main goal of our experiment is to illustrate the effectiveness of universum over

```

-----
Layer (type)           Output Shape           Param #
=====
Conv2d-1               [1, 32, 64, 64]      2,400
BatchNorm2d-2          [1, 32, 64, 64]      0
MaxPool2d-3            [1, 32, 32, 32]      0
Conv2d-4               [1, 64, 32, 32]      51,200
BatchNorm2d-5          [1, 64, 32, 32]      0
MaxPool2d-6            [1, 64, 16, 16]      0
Conv2d-7               [1, 128, 16, 16]     204,800
BatchNorm2d-8          [1, 128, 16, 16]     0
MaxPool2d-9            [1, 128, 8, 8]       0
Linear-10              [1, 128]              1,048,576
Linear-11              [1, 64]               8,192
Linear-12              [1, 1]                64
=====

```

Fig. 6 Network used for MVTec-AD results in Table 5

inductive learning for one class problems. Hence, we stick to a simple LeNet architecture shown in Fig. 6.

Finally, for both the above architectures we use bias = False for convolution operations and set $\epsilon = 10^{-4}$, Affine = False for BatchNorm. Additionally, we use a leaky ReLU activation after every max-pool operation.

B.2 Model parameters for Table 1

The optimal model parameters for the tabular datasets is provided in Table 7.

Table 7 Optimal model parameters for the tabular data sets

Class	λ	Learning rate	C_U/C	Epoch	Batch size
Thyroid	10^{-6}	10^{-3}	5.0	500	100
Arrhythmia	0.01	10^{-3}	0.001	300	100
Abalone	0.1	10^{-3}	0.01	300	100

B.3 Model parameters for Table 2 (CIFAR-10)

B.3.1 DOC and DOC³ model parameters used in Table 2

There are several hyper-parameters to be tuned for DOC and DOC³. To simplify our analysis we fix a few of these parameters following prior research.

- Unlike previous works like (Ruff et al., 2018; Goyal et al., 2020), we uniformly use an SGD optimizer with batch_size = 256. Although, training for each class represent a completely different problem, we adopt this to maintain consistency and isolate out the effect of optimizers for DOC versus DOC³ performances.

Table 8 Optimal model parameters for CIFAR-10 results for DOC

Class	λ	SGD (Learning rate)
Airplane	0.5	0.005
Automobile	1.0	0.001
Bird	0.5	0.005
Cat	1.0	0.001
Deer	1.0	0.001
Dog	0.5	0.001
Frog	0.5	0.001
Horse	1.0	0.001
Ship	0.5	0.001
Truck	0.001	0.001

Table 9 Optimal model parameters for CIFAR-10 results for DOC³

Class	λ	SGD (Learning rate)	C_U/C
Airplane	0.1	0.005	0.5
Automobile	0.05	0.0005	0.1
Bird	0.05	0.005	0.5
Cat	0.1	0.005	1.0
Deer	0.1	0.001	1.0
Dog	0.05	0.001	1.0
Frog	0.1	0.005	0.5
Horse	0.05	0.0005	1.0
Ship	0.05	0.005	0.1
Truck	0.05	0.0001	0.05

- For DOC we fix the total number of iterations for gradient updates to 300. Except for class ‘DOG’ and ‘Truck’ we use 400 and 50 respectively. For DOC³ we fix it to 350. This is in the same range as Ruff et al. (2018), and hence incurs similar computation complexity as the baseline DOCC and DROCC algorithms.
- Finally for DOC³ we fix $\Delta = 0$.

With the above hyper parameters fixed our best selected remaining hyper parameters for DOC and DOC³ are provided in Tables 8 and 9 respectively.

B.3.2 DOC (DA/OE) model parameters in Table 2

Next, we provide the optimal model parameters for the DOC (DA/OE) setting in Table 10. For the DOC (DA/OE) following Goyal et al. (2020) we introduce the universum samples as negative class in a standard binary hinge loss. The explicit form of this loss is also discussed in “Appendix C.1.1” in Eq. (C1). Here we set $C^+ = C^- = 1$.

B.3.3 Model parameters for DROCC-LF under OE versus Universum setting

For the DROCC-LF (OE) we use the same implementation as in Goyal et al. (2020). For the DROCC-LF under universum setting we replace the binary cross entropy loss used in Goyal

Table 10 DOC (DA/OE) parameters CIFAR-10

CIFAR-10 (Class)	$\lambda = \frac{1}{2C}$	SGD learning rate
Airplane	1.0	5×10^{-3}
Automobile	0.01	5×10^{-4}
Bird	0.5	5×10^{-4}
Cat	1.0	10^{-4}
Deer	1.0	10^{-4}
Dog	0.5	10^{-3}
Frog	0.5	5×10^{-3}
Horse	0.01	10^{-3}
Ship	1.0	10^{-3}
Truck	0.01	10^{-4}

Table 11 DROCC (OE) and DROCC (Univ) optimal (radius) parameters used in Table 2

CIFAR-10 (Class)	DROCC-LF (OE)	DROCC-LF (Univ)
Airplane	8	8
Automobile	32	8
Bird	16	8
Cat	16	16
Deer	16	16
Dog	16	8
Frog	32	32
Horse	16	32
Ship	8	8
Truck	8	8

et al. (2020) with the universum loss in (6) (see Algo. 1). Here we use the same notations as also used in Goyal et al. (2020). Further, as in Sect. 5.3.1 we replace the relu operator $[x]_+$ with the softplus operator for the loss functions.

We adopt the same LeNet architecture used in Ruff et al. (2018), Goyal et al. (2020) (see Fig. 4). Finally, we run the experiments over 10 runs and report the best AUC over the range of parameters recommended in Goyal et al. (2020) (Sect. 5). That is learning rate = 10^{-4} , radius (r) in range of $\sqrt{d} = \{8.0, 16.0, 32.0\}$. Here, for both the methods we use Adam and fix the number of ascent steps = 10 and batch size = 256 and total epochs = 350. The remaining parameters are set to default values. The final optimal parameters selected for the different classes is provided in Table 11.

Caveat(s): We found a few caveats while running the DROCC-LF experiments. One major caveat is that the gradient ascent steps are prone to instabilities. Note that the DROCC-LF algorithm (Algo. 2 in Goyal et al. (2020)) scales the perturbation direction (h) by the norm of the gradient vector. This results to severe gradient explosion. Appropriate measures to alleviate this issue has to be taken. Another major caveat is that the additional gradient ascent updates results to high computation complexity. For example, for the experiments presented in this paper a typical DROCC-LF run (350 epoch) takes $\sim 10^4$ secs compared to $\sim 10^3$ sec without the adversarial updates. The system configuration used here is,

Table 12 Optimal model parameters for F-MNIST results for DOC

Class	λ	SGD (learning rate)
T-shirt	1×10^{-5}	1×10^{-4}
Trouser	2.0	1×10^{-4}
Pullover	2.0	1×10^{-4}
Dress	1.0	1×10^{-4}
Coat	0.1	1×10^{-4}
Sandal	0.1	0.001
Shirt	1.0	1×10^{-4}
Sneaker	1×10^{-5}	0.001
Bag	2.0	2.5×10^{-4}
Ankle boot	0.001	1×10^{-4}

Table 13 Optimal model parameters for F-MNIST results for DOC³

Class	λ	SGD (learning rate)	C_U/C
T-shirt	1×10^{-4}	0.005	2.0
Trouser	1×10^{-5}	0.005	2.0
Pullover	1×10^{-8}	0.005	1.0
Dress	0.01	0.005	2.0
Coat	0.01	0.005	2.0
Sandal	1×10^{-4}	0.001	1.0
Shirt	1×10^{-8}	0.005	2.0
Sneaker	1×10^{-4}	0.009	2.0
Bag	0.0	0.001	5.0
Ankle boot	0.01	0.005	2.0

- CPU = AMD Ryzen 9 5950X 16 Core.
- RAM = 32 GB.
- GPU = NVIDIA GeForce RTX 3080
- CUDA = 11.4

B.4 Model parameters for Table 3 (FMNIST)

B.4.1 DOC and DOC³ model parameters used in Table 3

The optimal model parameters for DOC, DOC³ and DOC(DA/OE) are provided in Tables 12, 13, 14 respectively.

B.4.2 Model parameters for DROCC-LF under OE versus Universum setting for F-MNIST data

The optimal model parameters are provided in Table 15.

B.5 Model parameters for Table 5 (MVTec-AD)

Table 14 DOC (DA/OE) parameters F-MNIST

F-MNIST (Class)	$\lambda = \frac{1}{2C}$	SGD learning rate
T-shirt	0.001	1×10^{-4}
Trouser	0.001	1×10^{-4}
Pullover	1×10^{-5}	0.001
Dress	0.1	1×10^{-4}
Coat	1×10^{-9}	0.001
Sandal	0.1	0.001
Shirt	1×10^{-9}	1×10^{-4}
Sneaker	0.001	1×10^{-4}
Bag	0.1	1×10^{-4}
Ankle boot	0.001	0.005

Table 15 DROCC-LF (OE) and DROCC-LF (Univ) optimal (radius) parameters used in Table 3

F-MNIST (Class)	Radius DROCC-LF (OE)	DROCC-LF (Univ)	λ DROCC-LF (OE)	DROCC-LF (Univ)	C_U DROCC-LF (Univ)
T-shirt	8	8	0.1	0.1	2.0
Trouser	2	24	1.0	0.1	1.0
Pullover	8	16	0.1	0.1	0.1
Dress	8	8	0.1	0.1	1.0
Coat	16	8	1.0	0.01	2.0
Sandal	2	8	1.0	1.0	0.1
Shirt	2	16	0.1	1.0	2.0
Sneaker	16	16	0.1	1.0	0.1
Bag	2	8	1.0	0.01	0.1
Ankle boot	2	24	0.1	0.1	2.0

Here we provide the optimal model parameters selected and used to reproduce the DOC and DOC³ results in Tables 5 and 6. For this set of experiments we use the Adam optimizer with `batch_size = 100`. Further, to simplify model selection we fix the total number of iterations to 1000, and $\Delta = 0$. The optimal model parameters for DOC and DOC³ is provided in Table 16. Finally we also provide the optimal hyperparameters for the DOC (DA/OE) algorithm in Table 17.

Appendix C Additional experiments and results

C.1 Comparisons of disjoint auxiliary (or outlier exposure) versus Universum settings

In this section we highlight the differences between the universum versus the ‘Disjoint Auxiliary data’ setting used in Dhar (2014) (see Sect. 4.3) and Hendrycks et al. (2018), Goyal et al. (2020) etc. As discussed in the Sect. 4 a major difference is the assumption that the


Table 16 Optimal model parameters for MVTEC-AD results

Method		λ	Adam (learning rate)	C_U/C
Leather				
DOC		0.01	10^{-5}	–
DOC ³	Noise	0.01	10^{-5}	0.01
	Obj	0.01	10^{-5}	0.1
	Text	0.01	10^{-6}	0.01
Wood				
DOC		0.005	10^{-4}	–
DOC ³	Noise	0.005	5×10^{-6}	1.0
	Obj	0.005	10^{-5}	1.0
	Text	0.05	10^{-5}	0.1
Tile				
DOC		0.01	10^{-5}	–
DOC ³	Noise	0.1	5×10^{-6}	0.1
	Obj	0.005	10^{-4}	2.0
	Text	0.1	5×10^{-6}	0.1
Carpet				
DOC		1.0	5×10^{-4}	–
DOC ³	Noise	0.001	10^{-5}	0.01
	Obj	0.005	5×10^{-5}	2.0
	Text	1.0	10^{-3}	10^{-5}

Table 17 DOC (DA/OE) model parameters for MVTEC results

Normal	Univ. Anomaly	$\lambda = \frac{1}{2C}$	Learning rate
Leather	Noise	0.001	10^{-6}
	Obj	0.01	10^{-5}
	Text	0.0001	10^{-5}
Wood	Noise	0.01	10^{-5}
	Obj	0.01	10^{-6}
	Text	0.01	10^{-6}
Tile	Noise	0.005	10^{-6}
	Obj	0.001	10^{-4}
	Text	0.1	10^{-6}
Carpet	Noise	0.001	10^{-6}
	Obj	0.01	10^{-5}
	Text	0.005	10^{-5}

Fig. 7 Decision boundaries for *one class SVM* versus one class U-SVM versus binary SVM. Typical parameters, *one class SVM* ($C = 5$), one class U-SVM ($C_U = 10^{-3}$, $\Delta = 0$), binary SVM ($C = 10$). TP = True Positive, TN = True Negatives



universum samples act as contradictions to the unseen anomalous class (see Definition (2)). Methods using the ‘Disjoint Auxiliary’ setting do not use this assumption and formulate a loss function which only contradicts the ‘normal’ class. Such approaches have also been called ‘Supervised OE’ in Ruff et al. (2021) or ‘Limited Negatives’ in Goyal et al. (2020). In this section we take a more pedagogical approach to highlight the differences between Universum versus ‘Disjoint Auxiliary’ setting. For simplicity we use a binary classifier as an exemplar of this ‘Disjoint Auxiliary’ setting. That is, we build a binary classifier with ‘+1’ (normal samples) and ‘-1’ (contradiction a.k.a universum) samples. Note that, such an approach is philosophically inconsistent following Def. 2; where the universum samples are assumed to not follow the same distribution as both the normal (‘+1’) and anomalous (‘-1’) class. Using the universum samples as (‘-1’) class violates the assumption that universum follows a different distribution than the anomalous class. To further confirm our theoretical analysis we provide a simple synthetic example in example in C.1.1.

C.1.1 Synthetic experiment

For our synthetic example, we use synthetic data generated using normal distribution $\mathcal{N}(\mu, \sigma)$. For illustration we use,

- Normal Class (+1): $\mu = (1.0, 1.0)$, $\sigma = (0.25, 1.0)$.
- Anomaly Class (-1): $\mu = (0.25, 1.0)$, $\sigma = (0.25, 1.0)$.
- Contradictions: $\mu = (0.75, 6.0)$, $\sigma = (0.25, 1.0)$.

Additionally we use,

- No. of Training samples (+1 class) = 10
- No. of Test samples (+1, -1) class = 1000 per class.
- No. of Universum samples = 1000.

Note, that in the above synthetic example the discriminative power is mostly contained in the 1st dimension. Having ‘good’ universum samples can incorporate this additional information by contradicting the 2nd dimension while estimating the decision rule. This is also seen from Fig. 7. Figure 7 provides the decision boundaries obtained under inductive (3) versus universum settings (6) using only linear parameterization. Under linear parameterization the formulations reduces to standard SVM formulations so we refer them as one class SVM and one class U-SVM respectively. Finally, we also provide the decision boundary using a binary SVM, which serves as a representative for DA/OE-extension (Cherkassky & Mulier, 2007).

$$\min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w}\|_2^2 + C \sum_{i=1}^n \left[C^+ [1 - \mathbf{w}^\top \phi(\mathbf{x}_i)]_+ + C^- [1 + \mathbf{w}^\top \phi(\mathbf{x}_i)]_+ \right] \quad (\text{C1})$$

where, $[x]_+ = \max(0, x)$. For the binary SVM we use the universum as (-1) class, and adopt cost-sensitive formulation with a cost ratio $\frac{C^+}{C^-} = \frac{\#_{\text{univ}}}{\#_{\text{train}}} = \frac{1000}{10}$, to handle the class imbalance.

As seen from Fig. 7, using binary formulation in this universum setting does not correctly capture information available through the contradiction samples. That is, discriminating between normal and contradiction samples does not provide a good classifier for normal versus anomaly classification. The one class SVM although correctly classifies the positive samples (TP = 100%); does not perform good on future test samples. Using the universum samples, we can incorporate the additional information that the decision boundary should align along the vertical axis to have maximal contradiction (following Prop. 2). And by doing so, it improves the test performance over the one class SVM solution.

C.1.2 Ablation study hyperparameters

The DOC³ algorithm mainly introduces two additional hyper-parameters C_U and Δ compared to its inductive counterpart. The success of such an advanced technique depends on careful tuning of the hyperparameters. In this section we perform an ablation study of the $\frac{C_U}{C}$ and the Δ hyperparameters. To simplify we present the results for the CIFAR-10 data. Analysis using F-MNIST and MVTec-AD data provides similar conclusions.

Figures 8 and 9 provides the average \pm std. deviation of the AUC values over 10 experiment runs for varying $\frac{C_U}{C}$ —ratios and Δ values respectively. The experiment follows the same setting as in Sect. 5.3.1. Further all the other model parameters are set to their optimal values reported in Table 16. As seen from the figures, the model performance significantly varies for different $\frac{C_U}{C}$ -values (specifically for automobile, deer, dog, frog etc.). On the contrary, the DOC³ model performance seems relatively stable for varying Δ values (see Fig. 9). Such behavior is also seen for the other datasets. In line with this analysis throughout the paper we fix $\Delta = 0$ and follow the current norm of reporting the best model’s results over a small subset of hyperparameters. But this is far from practical. This motivates advanced mechanisms for optimal selection of this hyper parameter, which is still an open research topic. From our prior experiments, we found C_U/C in the range of [0.01, 2.0] provides reasonable performance in practice.

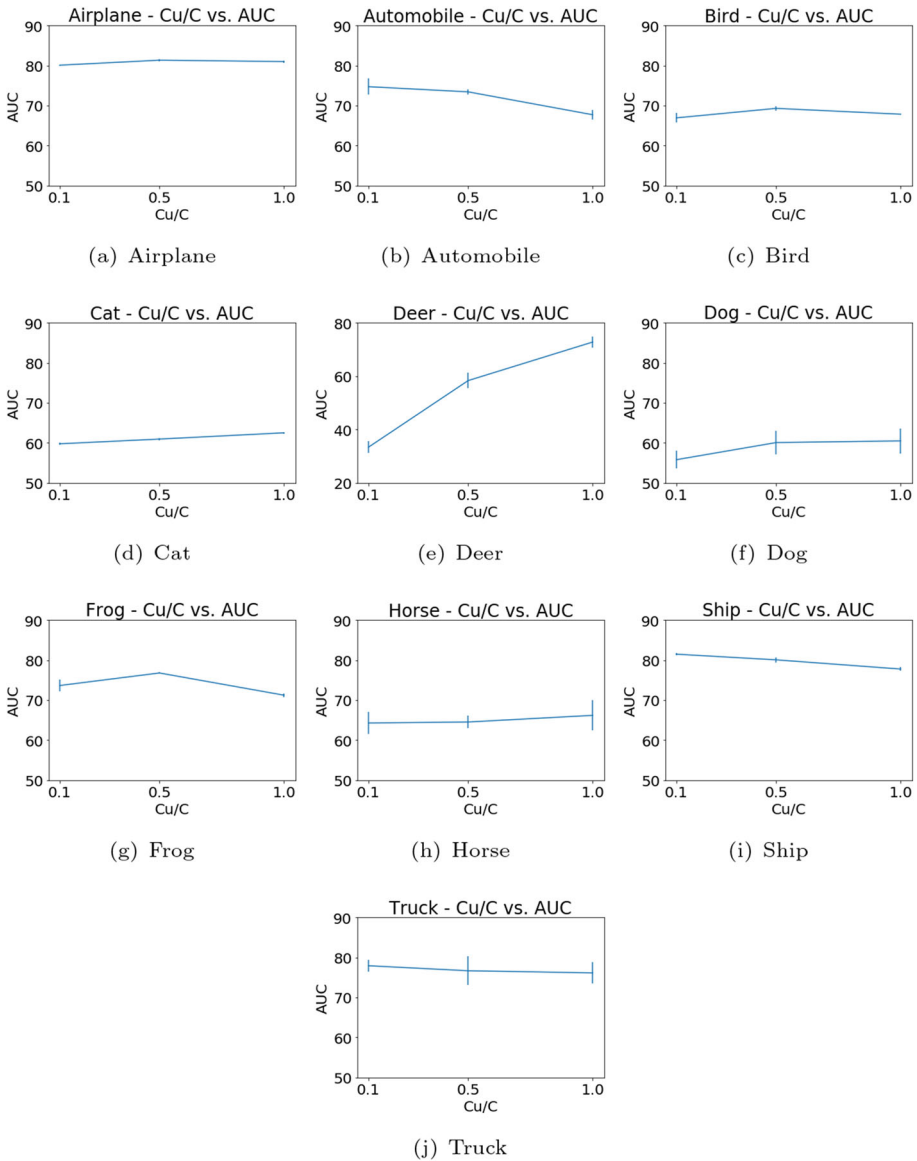


Fig. 8 Ablation study—AUC values for varying C_U/C values

C.2 Re-run of deep-SVDD (Ruff et al., 2018) and DROCC (Goyal et al., 2020) CIFAR-10 Results

C.2.1 Deep one class classification deep-SVDD results

For the Deep-SVDD results we see very similar results for our run except the ‘Frog’ and ‘Dog’ classes (see Table 18 and 19); where the difference are not too significant. Hence, we report the results as presented in the paper.

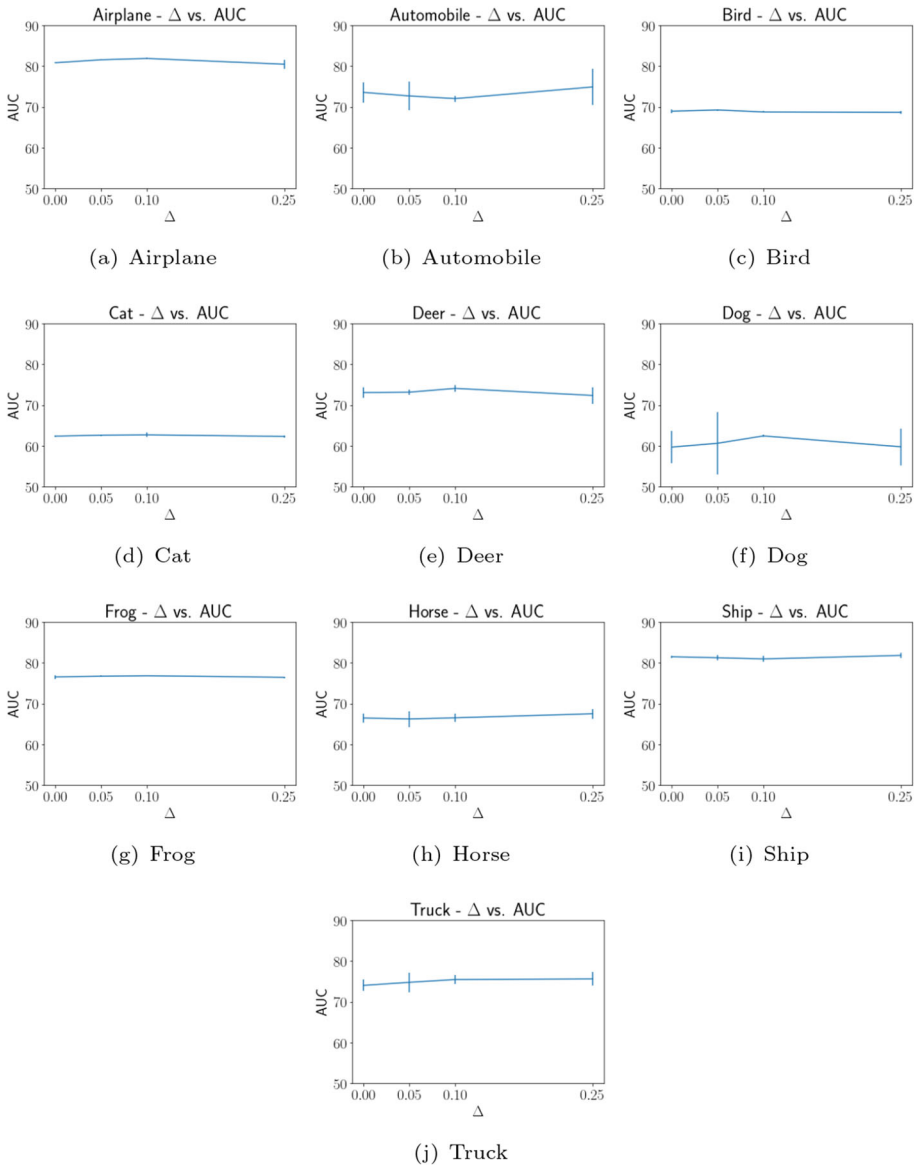


Fig. 9 Ablation study—AUC values for varying Δ values

C.2.2 Deep robust one class classification (DROCC) results

In Table 20, we report the results of our run with two different scaling. For the ‘all-class’ scale we use the scale used in the original DROCC paper (Goyal et al., 2020) i.e. $\mu = (0.4914, 0.4822, 0.4465)$ and standard deviation, $\sigma = (0.247, 0.243, 0.261)$. Note that, this scale is computed using the pixel values for all the classes. This in general is not available during training a one class classifier. Alternatively, ‘no-prior’ scale also reports the results

Table 18 Reproducing Deep-SVDD results in Ruff et al. (2018)

Class	Ruff et al. (2018)	Deep-SVDD (our Run)	
		$\nu = 0.1$	$\nu = 0.01$
Airplane	61.7 ± 4.1	61.0 ± 1.6	61.1 ± 1.4
Automobile	65.9 ± 2.1	60.4 ± 1.7	60.4 ± 1.7
Bird	50.8 ± 0.8	48.6 ± 0.6	48.6 ± 0.6
Cat	59.1 ± 1.4	57.7 ± 0.9	57.8 ± 1.0
Deer	60.9 ± 1.1	56.2 ± 0.7	56.3 ± 0.8
Dog	65.7 ± 2.5	63.4 ± 1.1	63.5 ± 1.3
Frog	67.7 ± 2.6	56.8 ± 1.8	56.9 ± 2.0
Horse	67.3 ± 0.9	59.9 ± 1.9	59.8 ± 1.9
Ship	75.9 ± 1.2	77.1 ± 1.0	76.9 ± 1.0
Truck	73.1 ± 1.2	66.9 ± 0.6	66.9 ± 0.7

Table 19 Reproducing Deep-SVDD (soft-boundary) results in Ruff et al. (2018)

Class	Ruff et al. (2018)	Deep-SVDD soft-boundary (Our Run)	
		$\nu = 0.1$	$\nu = 0.01$
Airplane	61.7 ± 4.1	61.9 ± 1.6	62.5 ± 1.9
Automobile	65.9 ± 2.1	61.6 ± 1.7	62.6 ± 2.0
Bird	50.8 ± 0.8	48.0 ± 0.9	45.9 ± 1.6
Cat	59.1 ± 1.4	56.6 ± 1.3	56.0 ± 1.8
Deer	60.9 ± 1.1	56.2 ± 0.8	56.1 ± 1.2
Dog	65.7 ± 2.5	61.9 ± 0.9	60.7 ± 1.7
Frog	67.7 ± 2.6	59.8 ± 1.8	61.0 ± 1.5
Horse	67.3 ± 0.9	61.5 ± 1.6	61.3 ± 1.5
Ship	75.9 ± 1.2	77.7 ± 0.9	76.7 ± 0.8
Truck	73.1 ± 1.2	67.5 ± 0.9	68.7 ± 1.4

using a scale using $\mu = (0.5, 0.5, 0.5)$ and standard deviation, $\sigma = (0.5, 0.5, 0.5)$. This scale does not need additional information from the other class's pixel values. We do not see a significant difference using these different scales. Although our re-runs show a significant difference for the 'ship' class between our results and the paper. We report the results of our re-run using the 'no-prior' scale in Table 2.

C.3 Re-run of PLAD (Cai & Fan, 2022) results on F-MNIST data

This section provides a detailed discussion of our re-run of the PLAD experiments on FMNIST data. PLAD (Cai & Fan, 2022) is a recent adversarial based work which improves upon DROCC (Goyal et al., 2020). We found several caveats and discrepancies in the PLAD implementation <https://openreview.net/forum?id=-Xdts90bWZ3>. This is highlighted below,

1. The network architecture discussed in paper and the implementation are different. The implementation does not use any non linearity (relu) operators for the final linear layers. Without the non-linear operators a cascade of affine operations can simply be replaced

Table 20 Reproducing DROCC results in Goyal et al. (2020)

Class	Goyal et al. (2020)	DROCC (all-class scale)	DROCC (no-prior scale)
Airplane	81.66 ± 0.22	79.99 ± 1.65	79.24 ± 1.95
Automobile	76.73 ± 0.99	74.61 ± 2.57	74.92 ± 2.66
Bird	66.66 ± 0.96	69.56 ± 0.94	68.29 ± 1.53
Cat	67.13 ± 1.51	54.54 ± 3.71	62.25 ± 2.67
Deer	73.62 ± 2.00	65.85 ± 2.94	70.34 ± 2.68
Dog	74.43 ± 1.95	66.47 ± 3.16	66.18 ± 2.09
Frog	74.42 ± 0.92	70.64 ± 2.40	68.16 ± 2.12
Horse	71.39 ± 0.22	70.18 ± 2.42	71.33 ± 4.57
Ship	80.01 ± 1.69	63.58 ± 7.88	62.39 ± 10.33
Truck	76.21 ± 0.67	75.12 ± 1.92	76.58 ± 1.94

by a single affine operation. For our implementation we correct it and use non-linear relu operators in-between the linear operations.

2. The classes are differently scaled. Class ‘T-shirt’ and ‘Trouser’ uses a generic mean = [0.5], std = [0.5] scale, while the other classes are scaled to their respective mean and std. Further, as previously discussed for DROCC re-runs a class specific normalization is not practical for one-class problems as we cannot expect the knowledge of class labels apriori during training or testing. Throughout our paper we use a generic mean = [0.5], std = [0.5] scale.
3. Finally the results provided in Cai and Fan (2022) is the best AUC obtained during the entire training process (and not at the final stopping criteria). While such an approach may still be practical for supervised learning, where intermediate validation AUCs may provide guidance to select the optimal model. It is impossible to select an intermediate best model during the training process for one-class (unsupervised) learning problems. For this work Table 3 provides the avg ± std. of the test AUCs obtained at the final step of training. For interested readers Table 21 also provides the avg ± std. of the best test AUCs obtained during the training process as used in Cai and Fan (2022).

Table 21 Avg \pm standard deviation of the best AUCs during training process for one-vs-rest anomaly detection on F-MNIST

F-MNIST	T-shirt	Trouser	Pullover	Dress	Coat	Sandal	Shirt	Sneaker	Bag	Ankle boot
OC-SVM [†]	86.1	93.9	85.6	85.9	84.6	81.3	78.6	97.6	79.5	97.8
IF [†]	91.0	97.8	87.2	93.2	90.5	93.0	80.2	98.2	88.7	95.4
DSVDD [†]	79.1	94.0	83.0	82.9	87.0	80.3	74.9	94.2	79.1	93.2
DROCC [†]	88.1	97.7	87.6	87.7	87.2	91.0	77.1	95.3	82.7	95.9
PLAD [‡]	90.35 \pm 1.1	97.7 \pm 1.1	89.3 \pm 1.3	91.2 \pm 0.4	92.2 \pm 0.5	96.1 \pm 0.9	81.8 \pm 0.9	97.8 \pm 0.2	86.9 \pm 1.3	99.0 \pm 0.1
DOC (Hinge Eq. 3)	74.7 \pm 7.7	90.9 \pm 0.2	77.0 \pm 3.7	85.4 \pm 2.0	81.0 \pm 2.3	96.1 \pm 0.7	66.8 \pm 1.3	92.1 \pm 4.1	77.3 \pm 1.2	94.3 \pm 2.1
DOC (DA/OE)	86.6 \pm 3.4	97.0 \pm 0.3	83.1 \pm 2.0	89.6 \pm 1.2	88.7 \pm 1.6	95.2 \pm 0.7	72.5 \pm 3.2	97.6 \pm 0.6	80.9 \pm 6.0	98.8 \pm 0.1
DOC ³ (univ)	88.3 \pm 2.2	97.4 \pm 0.4	85.3 \pm 2.2	91.6 \pm 0.6	90.8 \pm 1.1	96.0 \pm 2.3	75.7 \pm 1.3	98.2 \pm 0.2	80.1 \pm 2.3	98.8 \pm 0.2
DROCC-LF(OE)	91.6 \pm 1.0	97.0 \pm 0.2	88.1 \pm 0.9	91.4 \pm 0.4	91.4 \pm 0.2	96.2 \pm 1.0	82.4 \pm 0.6	97.9 \pm 0.1	87.3 \pm 1.5	99.1 \pm 0.1
DROCC-LF(univ)	93.8 \pm 0.3	98.3 \pm 0.1	89.6 \pm 0.08	92.8 \pm 0.7	92.3 \pm 0.1	96.1 \pm 0.7	82.1 \pm 0.2	98.3 \pm 0.1	88.2 \pm 0.3	99.3 \pm 0.05

[†]Results reported in Cai and Fan (2022). [‡]Our re-run of the algorithm in Cai and Fan (2022)

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., & Isard, M. (2016). Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI})* (Vol. 16, pp. 265–283).
- Akçay, S., Atapour-Abarghouei, A., & Breckon, T. P. (2018). Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Asian conference on computer vision* (pp. 622–637). Springer.
- Bartlett, P. L., & Mendelson, S. (2002). Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3, 463–482.
- Bergman, L., & Hoshen, Y. (2020). Classification-based anomaly detection for general data. arXiv preprint [arXiv:2005.02359](https://arxiv.org/abs/2005.02359)
- Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2019). Mvtec ad—A comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 9592–9600).
- Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. (2000). Lof: Identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on management of data* (pp. 93–104).
- Cai, J., & Fan, J. (2022). Perturbation learning based anomaly detection. In A.H. Oh, A. Agarwal, D. Belgrave, & K. Cho (Eds.), *Advances in neural information processing systems*. <https://openreview.net/forum?id=Xdts90bWZ3>
- Caron, M., Bojanowski, P., Joulin, A., & Douze, M. (2018). Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 132–149).
- Carrara, F., Amato, G., Brombin, L., Falchi, F., & Gennaro, C. (2020). Combining gans and autoencoders for efficient anomaly detection. arXiv preprint [arXiv:2011.08102](https://arxiv.org/abs/2011.08102)
- Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. arXiv preprint [arXiv:1901.03407](https://arxiv.org/abs/1901.03407)
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3), 1–58.
- Chang, C.-C., & Lin, C.-J. (2001). Training v-support vector classifiers: Theory and algorithms. *Neural Computation*, 13(9), 2119–2147.
- Chapelle, O., Agarwal, A., Sinz, F. H., & Schölkopf, B. (2008). An analysis of inference with the universum. In *Advances in neural information processing systems* (pp. 1369–1376).
- Chen, S., & Zhang, C. (2009). Selecting informative universum sample for semi-supervised learning. In *IJCAI* (pp. 1016–1021).
- Cherkassky, V., Dhar, S., & Dai, W. (2011). Practical conditions for effectiveness of the universum learning. *IEEE Transactions on Neural Networks*, 22(8), 1241–1255.
- Cherkassky, V., & Mulier, F. M. (2007). *Learning from data: Concepts, theory, and methods*. Wiley.
- Cortes, C., Gonzalvo, X., Kuznetsov, V., Mohri, M., & Yang, S. (2017). Adanet: Adaptive structural learning of artificial neural networks. In *International conference on machine learning* (pp. 874–883). PMLR.
- Das, S., Islam, M. R., Jayakodi, N. K., & Doppa, J. R. (2018). Active anomaly detection via ensembles. arXiv preprint [arXiv:1809.06477](https://arxiv.org/abs/1809.06477)
- Dhar, S. (2014). Analysis and extensions of universum learning.
- Dhar, S., & Cherkassky, V. (2017). Universum learning for svm regression. In *2017 International joint conference on neural networks (IJCNN)* (pp. 3641–3648). IEEE.
- Dhar, S., Cherkassky, V., & Shah, M. (2019). Multiclass learning from contradictions. In *Advances in neural information processing systems* (pp. 8400–8410).
- Dhar, S., & Cherkassky, V. (2015). Development and evaluation of cost-sensitive universum-SVM. *IEEE Transactions on Cybernetics*, 45(4), 806–818.
- Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1). MIT Press.
- Goyal, S., Raghunathan, A., Jain, M., Simhadri, H. V., & Jain, P. (2020). Drocc: Deep robust one-class classification. In *International conference on machine learning* (pp. 3711–3721). PMLR.
- Hendrycks, D., Mazeika, M., & Dietterich, T. (2018). Deep anomaly detection with outlier exposure. arXiv preprint [arXiv:1812.04606](https://arxiv.org/abs/1812.04606)
- Hoffmann, H. (2007). Kernel PCA for novelty detection. *Pattern Recognition*, 40(3), 863–874.
- Huang, C., Cao, J., Ye, F., Li, M., Zhang, Y., & Lu, C. (2019). Inverse-transform autoencoder for anomaly detection. arXiv preprint [arXiv:1911.10676](https://arxiv.org/abs/1911.10676)
- Huang, C., Ye, F., Cao, J., Li, M., Zhang, Y., & Lu, C. (2019). Attribute restoration framework for anomaly detection. arXiv preprint [arXiv:1911.10676](https://arxiv.org/abs/1911.10676)
- Huang, C., Ye, F., Zhang, Y., Wang, Y.-F., & Tian, Q.: Esad: End-to-end deep semi-supervised anomaly detection. arXiv preprint [arXiv:2012.04905](https://arxiv.org/abs/2012.04905)

- Huang, S.-H., & Pan, Y.-C. (2015). Automated visual inspection in the semiconductor industry: A survey. *Computers in Industry*, 66, 1–10.
- Jolliffe, I. T. (2002). *Principal component analysis*. Springer.
- Knorr, E. M., Ng, R. T., & Tucakov, V. (2000). Distance-based outliers: Algorithms and applications. *The VLDB Journal*, 8(3), 237–253.
- Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *2008 Eighth IEEE international conference on data mining* (pp. 413–422). IEEE.
- Makhzani, A., & Frey, B. (2014). Winner-take-all autoencoders. arXiv preprint [arXiv:1409.2752](https://arxiv.org/abs/1409.2752)
- Masci, J., Meier, U., Cireşan, D., & Schmidhuber, J. (2011). Stacked convolutional auto-encoders for hierarchical feature extraction. In *International conference on artificial neural networks* (pp. 52–59). Springer.
- Massoli, F. V., Falchi, F., Kantarci, A., Akti, Ş., Ekenel, H. K., & Amato, G. (2020). Mocca: Multi-layer one-class classification for anomaly detection. arXiv preprint [arXiv:2012.12111](https://arxiv.org/abs/2012.12111)
- Neyshabur, B., Tomioka, R., & Srebro, N. (2015). Norm-based capacity control in neural networks. In *Conference on learning theory* (pp. 1376–1401).
- Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2020). Deep learning for anomaly detection: A review. arXiv preprint [arXiv:2007.02500](https://arxiv.org/abs/2007.02500)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 3, 2.
- Patel, R., & Toda, M. (1979). Trace inequalities involving Hermitian matrices. *Linear Algebra and its Applications*, 23, 13–20.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. (2011). Scikit-learn: Machine learning in python. *The Journal of machine Learning research*, 12, 2825–2830.
- Ramaswamy, S., Rastogi, R., & Shim, K. (2000). Efficient algorithms for mining outliers from large data sets. In *Proceedings of the 2000 ACM SIGMOD international conference on management of data* (pp. 427–438)
- Rosenberg, D. S., & Bartlett, P. L. (2007). The rademacher complexity of co-regularized kernel classes. In *Artificial intelligence and statistics* (pp. 396–403).
- Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. In *Proceedings of the IEEE*.
- Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., & Kloft, M. (2018). Deep one-class classification. In *International conference on machine learning* (pp. 4393–4402).
- Ruff, L., Vandermeulen, R. A., Goernitz, N., Binder, A., Müller, E., Müller, K.-R., & Kloft, M. (2019). Deep semi-supervised anomaly detection. arXiv preprint [arXiv:1906.02694](https://arxiv.org/abs/1906.02694)
- Schlegl, T., Seeböck, P., Waldstein, S. M., Schmidt-Erfurth, U., & Langs, G. (2017). Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International conference on information processing in medical imaging* (pp. 146–157). Springer.
- Schölkopf, B., Williamson, R., Smola, A., Shawe-Taylor, J., & Platt, J. (1999). Support vector method for novelty detection. In *Proceedings of the 12th international conference on neural information processing systems*.
- Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., & Williamson, R. C. (2001). Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7), 1443–1471.
- Schölkopf, B., Smola, A. J., Bach, F., et al. (2002). *Learning with Kernels: Support vector machines, regularization, optimization, and beyond*. MIT Press.
- Shawe-Taylor, J., Cristianini, N., et al. (2004). *Kernel methods for pattern analysis*. Cambridge University Press.
- Shen, C., Wang, P., Shen, F., & Wang, H. (2012). Uboost: Boosting with the universum. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(4), 825–832.
- Sinz, F., Chapelle, O., Agarwal, A., & Schölkopf, B. (2008). An analysis of inference with the universum. In *Advances in neural information processing systems* (Vol. 20, pp. 1369–1376). Curran.
- Sokolic, J., Giryas, R., Sapiro, G., & Rodrigues, M. R. (2016). Lessons from the rademacher complexity for deep learning.
- Tack, J., Mo, S., Jeong, J., & Shin, J. (2020). Csi: Novelty detection via contrastive learning on distributionally shifted instances. *Advances in Neural Information Processing Systems*, 33, 11839–11852.
- Tan, P.-N., Steinbach, M., & Kumar, V. (2016). *Introduction to data mining*. Pearson Education India.
- Tax, D. M., & Duin, R. P. (2004). Support vector data description. *Machine Learning*, 54(1), 45–66.
- Vapnik, V. (2006). *Estimation of dependences based on empirical data (information science and statistics)*. Springer.

- Weimer, D., Scholz-Reiter, B., & Shpitalni, M. (2016). Design of deep convolutional neural network architectures for automated feature extraction in industrial inspection. *CIRP Annals*, 65(1), 417–420.
- Weston, J., Collobert, R., Sinz, F., Bottou, L., & Vapnik, V. (2006). Inference with the universum. In *Proceedings of the 23rd international conference on machine learning* (pp. 1009–1016). ACM.
- Xiao, Y., Feng, J., & Liu, B. (2021). A new transductive learning method with universum data. *Applied Intelligence*, 1–13.
- Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms. arXiv preprint [arXiv:1708.07747](https://arxiv.org/abs/1708.07747)
- Zenati, H., Foo, C. S., Lecouat, B., Manek, G., & Chandrasekhar, V. R. (2018). Efficient gan-based anomaly detection. arXiv preprint [arXiv:1802.06222](https://arxiv.org/abs/1802.06222)
- Zhang, X., & LeCun, Y. (2017). Universum prescription: Regularization using unlabeled data. In *AAAI* (pp. 2907–2913).
- Zong, B., Song, Q., Min, M. R., Cheng, W., Lumezanu, C., Cho, D., & Chen, H. (2018). Deep autoencoding Gaussian mixture model for unsupervised anomaly detection. In *International conference on learning representations*. <https://openreview.net/forum?id=BJJLHbb0>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.