



A Multi-source Domain Adaption Approach to Minority Disk Failure Prediction

Wang Wang^{1,2}, Xuehai Tang^{1(✉)}, Biyu Zhou¹, Yangchen Dong¹,
Yuanhang Feng¹, Jizhong Han¹, and Songlin Hu¹

¹ Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
{wangwang, tangxuehai, zhoubiyu, dongyangchen, fengyuanhang,
hanjizhong, husonglin}@iie.ac.cn

² School of Cyber Security, University of Chinese Academy of Sciences,
Beijing, China

Abstract. Frequent happening of disk failures affects the reliability of the storage system, which can cause jittering of performance or even data loss of services and thus seriously threaten the quality of service. Although a host of machine (deep) learning-based disk failure prediction approaches have been proposed to prevent system breakdown due to unexpected disk failure, they are able to achieve high performance based on the assumption that the disk model has plenty of samples (especially failure samples). However, new disk models continuously appear in data centers with the evolution of disk manufacturing technology and the expansion of storage system capacity. Limited by the deploying time, these disk models have few failure samples and are called minority disks. The minority disks are widespread in large-scale data centers and contain amounts of disks while existing approaches cannot reach satisfying performance on such disks due to the lack of failure samples. What's worse, failure prediction models trained on other disk models cannot be directly applied to these minority disks either due to the commonly existing distribution shift among disk models. In this work, we propose DiskDA, a novel multi-source domain adaption-based solution that can fully utilize knowledge from other disk models to predict failures for minority disks having no failure samples. Our experimental results on real-world datasets show the superiority of DiskDA against previous approaches on minority disks with a few failure samples. What's more, DiskDA also shows its good adaptivity on minority disks having no failure samples, whereas previous works are unusable.

Keywords: Fault tolerance · Disk failure prediction · Domain adaption · Cloud computing

1 Introduction

Disk failures are common in modern large-scale data centers, accounting for more than 70% of hardware replacement events [5, 13, 16]. Frequent happening

of disk failure can lead to service performance jitter or even data loss which severely affects the availability and reliability of cloud applications [7, 17]. To ensure the availability and reliability of cloud applications from unexpected disk failures, operators should proactively predict the upcoming disk failure events before they actually happen, so as to take preventive measures in time, such as virtual machine migration.

The Self-Monitoring, Analysis, and Reporting Technology (SMART) has been widely implemented by hard disk drive (HDD) and solid-state drive (SSD) manufacturers to monitor the status of individual disk drives. The values of SMART attributes related to disk health status are helpful to disk health tendency assessment.

Recently, with the development of machine learning, a host of supervised learning-based approaches has been proposed to predict disk failures with the SMART values [10, 20, 22, 23]. With sufficient samples (both healthy and failure samples) provided, these methods are able to train binary classifiers and classify newly coming disk samples collected periodically from data centers to predict failures for each disk with high accuracy.

Table 1. Statistics of Disk Population

Data Center	Disk Type	Type Number	Type Percentage	Disk Number
Backblaze	Majority	12	11.65%	114,570
	Minority	91	88.35%	34,978
Tencent	Majority	8	13.33%	52,235
	Minority	52	86.67%	18,996

However, the condition of sufficient failure samples can hardly be satisfied by all disk models. With the evolution of disk manufacturing technology and the expansion of storage system capacity, disks from different models are continuously added to data centers. Limited by the deploying scale and time, the newly coming disk models usually have only a few or even no failure samples for most cases and are named minority disks [11, 24]. According to the studies [9, 24] in large-scale data centers (i.e. Backblaze, Tencent, and AliCloud), minority disks are generally existing in modern data centers. As shown in Table 1 [24], the minority disks dominate the disk models (over 85%) and contain a great number of disks (tens of thousands of disks). Unfortunately, the traditional supervised ML method cannot be applied to predict failures for minority disks, otherwise, it will suffer from over-fitting or cold-start issues [3, 6, 24, 25]. What’s worse, the prediction model trained on other disk models cannot be applied to minority disks either and we illustrate this in our extensive experiments. Because the commonly existing distribution shifts across disk models break the assumption of independent identical distribution holding between training and test set.

More recently, several transfer learning (TL)-based methods [2, 9, 15, 19, 24] and semi-supervised learning approaches [3, 6, 25] are proposed. Based on the fact that the failure modes are common for different disk models (e.g. all disks will fail due to too many bad sectors) [5], the TL-based methods try to adapt failure prediction knowledge extracted from other disk models to minority disk models by directly selecting samples similar to the minority disks as training set or transforming the SMART attribute distribution of minority disks to that of other disk models via a heuristic statistical model. However, the existing TL-based approaches can only transfer partial knowledge from a single source domain (other disk models) since they have to drop many useful samples due to their dissimilarity to minority disks or abandon critical features that hard to be transformed. Although using multiple source domains is more likely to introduce more failure modes, the large number of samples contained in it also means the complex distribution of source domains, which will lead to negative migration problems in the existing TL-based methods. In addition, all existing TL-based approaches need a certain number of failure samples of minority disks in their transfer procedure, which suggests they can only handle very limited cases. As for the semi-supervised learning approaches, they though can train their model with only healthy samples, large quantities of minority disk failure samples are needed to set appropriate classification thresholds.

In this work, we are exploring extracting the transferable failure modes from multiple source domains and aligning their semantics across source and target domains so that the full knowledge can be leveraged to enhance the failure prediction of minority disks. To this end, we model our problem as an unsupervised domain adaption problem and propose DiskDA, a multi-source domain adaption-based failure prediction approach for minority disks. It is able to extract failure modes from samples of multiple disk models and utilize them in minority disk (even with no failure samples) failure prediction with high performance. The goal is achieved because of two key designs. Firstly, although minority disks do not have complete class distribution, we find that the particularity of the distribution of disk samples can be leveraged to ensure the execution of domain adaption. Based on this, we use a representor to extract failure modes from source domain samples and align their semantics across two domains using only healthy samples in the target domain. And a Wasserstein distance measurement is adopted to guarantee the effectiveness of the domain adaption even when the distribution of two domains is distant. We also prove the rationality of this strategy by analyzing the generalization error bound in this case. Secondly, DiskDA adopts a confidence-based sample selection to filter out irrelevant samples in the source domains, so as to eliminate the negative transfer issue. By running the two processes alternatively, DiskDA can successfully extract transferable failure modes from multiple source domain disk models and utilize them in minority disk failure prediction with high accuracy.

The main contributions are summarized as:

- 1) We explore the problem of failure prediction for minority disk without failure samples so that it can be adaptive to all minority disk models.

- 2) To the best of our knowledge, we are the first to propose a Wasserstein distance-based domain adaption solution for the minority disk failure prediction problem and first analyze the generalization error bound theoretically under this condition.
- 3) Guided by the generalization error bound, we design a novel unsupervised domain adaption framework, DiskDA, to minimize the generalization error of the failure predictor in the target domain.
- 4) We conduct evaluations to demonstrate the superiority of DiskDA on 9 disk models from 3 vendors, collected from 2 large-scale data centers. The evaluation results reveal that DiskDA can improve the F1-score by an average of 20.02% compared with the best competitor when less than a dozen failure samples are provided. **More importantly, DiskDA can still obtain a satisfactory F1-Score of about 0.93 when no failure samples are provided (most minority disks face), while all existing TL-based approaches will fail.**

2 Related Work

- **Supervised Learning-Based Failure Prediction Approaches.** Li et al. [10] propose a Classification And Regression Trees (CART) based model which can give disks a health assessment. Xu et al. [20] present a Recurrent Neural Networks (RNN [12]) method to leverage sequential information in hard disk failure prediction. Yang et al. [22] design a disk failure prediction model by using L1-regularized logistic regression. Zhang et al. [23] adopt the Siamese network [4] to improve the applicability and adaptivity of the disk failure prediction model. All these supervised learning-based approaches can achieve high performance based on the assumption that large quantities of failure samples are provided. However, this is harsh for minority disks.
- **Semi-supervised Learning-Based Failure Prediction Approaches.** The main idea of the semi-supervised learning-based approach is to model the distribution of healthy samples and predict failure samples based on their reconstruction errors. Once the reconstruction errors surpass a predefined threshold, the disk samples are classified as failure samples. Jiang et al. [6] propose a GAN (Generative Adversarial Network)-based anomaly prediction approach that adopts an encoder-decoder-encoder architecture. They define the reconstruction error as the difference between two encoders' outputs and predict failures by comparing the error with a threshold. Zhou et al. [25] and Chakrabortii et al. [3] predict failures for SSDs with similar approaches. The performance of such approaches relies on the manually set thresholds and the operators are able to find appropriate thresholds only when a certain number of reconstruction errors of failure samples are provided. Since the failure samples of minority disks are limited, it is hard for such methods to reach satisfying performance in minority disk failure prediction.
- **Transfer Learning-Based Failure Prediction Approaches.** The target of the TL-based approach is to adapt a failure prediction model trained from

existing disk models (source domain) to the minority disk (target domain). MirelaMadalina Botezatu et al. [2] propose an instance-based de-bias approach. They select samples from the source domain disk model based on their similarity degree given by a domain classifier and a Regularized Greedy Forests (RGF [8]) trained on the augmented minority disk dataset is adopted as the failure prediction model. Xie et al. [19] select the source domain based on the performance similarity of the failure prediction model on the minority disk and each candidate source domain disk model. Then the minority disk failure prediction model is trained on the union set of source and target domains. Zhang et al. [24] propose to utilize the Kullback-Leibler divergence (KLD) of the specific SMART attribute to select the source domain disk model and adopt the Tradaboost algorithm trained on both domains as the minority disk failure prediction model. Sun et al. [15] take another approach and propose to use a statistic-based feature transformation to align cumulative SMART attribute (e.g. SMART_5 represents reallocated sector count) distribution. They find the same cumulative SMART attribute of disks from different vendors/models have similar distributions and align their distributions based on the ratio of failed to healthy devices so as to adapt the failure prediction model trained from one disk model to others. Lan et al. [9] also try to transfer knowledge from the source domain by utilizing a domain classifier to learn the domain invariant representation of source and target domain samples. It is worth noting that the domain invariant representation learning guided by the domain classifier will fail (gradient vanishing) if the distribution of the source and target domain is distant [18]. And this has been shown in their experiment where 50% of the transfer process (domain adaption) failed due to the large distribution divergence of source and target domain samples. In a word, existing TL-based can only utilize limited information from the source domain due to the drop of samples and critical attributes. In addition, they can only work when a certain number of failure samples from minority disks are provided, while this can be harsh for minority disks.

To sum up, DiskDA differs from previous approaches in three aspects:

- Compared to supervised learning-based approaches, DiskDA extracts failure prediction knowledge from large amounts of samples from other disk models. And this strategy protects DiskDA from overfitting caused by the limited failure samples of minority disks.
- DiskDA adopts a binary classifier built on labeled samples to automatically discriminate the healthy and failed samples rather than manually setting the classification threshold as semi-supervised learning approaches.
- Compared to existing TL-based approaches, DiskDA does not choose to drop samples or critical attributes but tries to fuse the distribution of source and target domain samples, so as to fully utilize the failure prediction knowledge from source domain disk models. DiskDA avoids the gradient vanishing problem by adopting the Wasserstein distance to guide the domain invariant representation learning process. Because the Wasserstein distance can always provide stable gradients no matter how distant the distributions are [18].

3 Motivation

3.1 Problem Statement

In the problem of minority disk failure prediction based on domain adaption (MDFP-DA), we suppose a labeled dataset $X^s = \{(x_i^s, y_i^s)_{i=1}^{n_s}\}$ including n_s samples from **multiple** disk models of the data center, which are sufficient to train a high precision prediction model. Furthermore, we assume a dataset $X^t = \{(x_m^t, y_m^t)_m\}$ from the minority disk, where x_i^t refers to the sample collected online in future and y_i^t is the corresponding label. The samples from X^s and X^t ①share the same feature space (this can be ensured by keeping their common SMART attributes), but ②follow different marginal distributions, \mathbb{P}_s and \mathbb{P}_t . Although X^t is unreachable in reality, we can collect quantities of healthy samples $X_H^t = \{(x_j^t, 0)_{j=1}^{n_t}\}$ from the minority disk through short-term deployment, which is always held in disk failure prediction [3, 6, 25]. And we denote the marginal distribution of the healthy samples (from both X^t and X_H^t) as \mathbb{P}_{t_H} . Here, we regard MDFP-DA as a binary classification problem and label the healthy samples as ‘0’, and the failure samples as ‘1’. Now we give the definition of the MDFP-DA problem:

Definition 1. *The MDFP-DA problem is to learn a transferable classification model $h(s)$ to minimize the risk $\epsilon_t(h) = Pr_{(x,y) \sim X^t}[h(x) \neq y]$ using X^s and X_H^t .*

3.2 Generalization Error Bound Analysis

We analyze the generalization error bound by introducing the unsupervised domain adaption problem. The unsupervised domain adaption studies the problem of adapting a classifier trained in the source domain to target domain ①sharing the same feature space while ②having different data distribution.

Obviously, in the case of given X^t , the MDFP-DA problem can be converted to an unsupervised domain adaption problem. Although the failure samples of the minority disk are lost in the MDFP-DA problem, we have amounts of its healthy samples X_H^t . Let $\epsilon_t(h)$ denote the generalization error bound of a classification function h in target domain t . Let $W_1(P, Q)$ denote the Wasserstein distance between P and Q . In this case, the following Theorem holds.

Theorem 1. *For any classification function h to the MDFP-DA problem satisfying K -Lipschitz, the following holds:*

$$\epsilon_t(h) \leq \epsilon_s(h) + 2KW_1(\mathbb{P}_s, \mathbb{P}_{t_H}) + \lambda + C \quad (1)$$

where C is the Wasserstein distance of distribution of target domain \mathbb{P}_t and its healthy samples \mathbb{P}_{t_H} .

Proof. See Appendix for details.

Considering the fact that the healthy samples dominate the whole disk samples (with an average ratio of 9997:10000 [2]), the distribution of samples of a disk model is actually similar to its healthy samples, which suggests that C is a small constant. We have verified this by randomly selecting 4 disk models and calculating C and $W_1(\mathbb{P}_{t_H}, \mathbb{P}_s)$. And the results show that C is small in scale of 10^{-3} and $W_1(\mathbb{P}_{t_H}, \mathbb{P}_s)$ are hundreds of times of C , so it can be ignored in practice.

Remark. Theorem 3.1 implies that the generalization error of a prediction model in the target domain (i.e., $\epsilon_t(h)$) is smaller than the sum of the generalization error of the prediction model in the source domain (i.e., $\epsilon_s(h)$), the Wasserstein distance of source domain samples and minority disk healthy samples (i.e., $W_1(\mathbb{P}_s, \mathbb{P}_{t_H})$), and a constant (i.e., $\lambda + C$) much smaller than the former two. In other words, the generalization error of the prediction model in the minority disk ($\epsilon_t(h)$) can be optimized if we are able to reduce $\epsilon_s(h)$ and $W_1(\mathbb{P}_s, \mathbb{P}_{t_H})$. Once $\epsilon_t(h)$ is optimized, the performance of the failure prediction model in the minority disk can be improved. To sum up, it not only proves the generalization error bound of the MDFP-DA problem but also indicates the optimization direction in the absence of failure samples.

4 Method

4.1 Overview of DiskDA

Figure 1 illustrates the framework of DiskDA. As seen, the DiskDA consists of two processes, ① the domain invariant representation learning and ② the confidence-based sample selection.

① The first process mainly involves three modules:

- **Representer:** a deep neural network that projects the samples from source domain disk models (X^s) and minority disk (X_H^t) into a unified latent space (representation vector in Fig. 1).
- **Distance Estimator:** measures the Wasserstein distance of representation vectors from X^s and X_H^t .
- **Failure Predictor:** classifies whether a sample from X^s is a failure sample based on their representation vectors.

The representer is able to reduce the Wasserstein distance of the source domain and minority disk samples (i.e. reducing $W_1(\mathbb{P}_s, \mathbb{P}_{t_H})$) under the guidance of distance estimator. In the meantime, it helps the failure predictor to reach high performance in the source domain (i.e., reducing $\epsilon_s(h)$) via extracting discriminant information. In this way, the performance of the failure predictor in the minority disk can be optimized according to Theorem 3.1.

② The second process mainly includes the sample selector module. Its main purpose is to avoid negative transfer which may occur in the first process. The principle is to eliminate the samples in X^s that hinder the further narrowing of the distance between X^s and X_H^t , which also corresponds to reducing $W_1(\mathbb{P}_s, \mathbb{P}_{t_H})$.

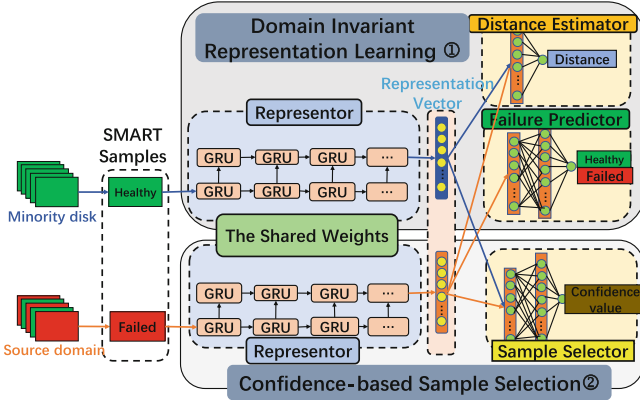


Fig. 1. The architecture of DiskDA.

In the training stage, DiskDA iterates the two processes alternatively. For example, the process ① runs every N (e.p., 100) iterations, and then the process ② runs M (e.p., 1) iterations. In this way, DiskDA can fully utilize the failure prediction knowledge of the source domain and lose the least information. The alternating iteration can be stopped until the parameters are converged or the iteration times reach a threshold.

In the online prediction, the SMART instances of a minority disk are collected daily to a sample pool. And these instances are combined to form samples as in the training stage. All these samples will be input into the repressor to generate corresponding representation vectors. Then the failure predictor will predict whether a sample is a failure sample based on its representation. Once a sample of a disk (minority disk) is predicted as a failure sample, it suggests that the disk will fail soon and the alarm system will inform the operator to repair/exchange the disk in time.

4.2 Domain Invariant Representation Learning

To effectively adapt the failure predictor trained in the source domain disks to the minority disk, we need the repressor to learn domain invariant representations of samples from both domains. That is, the distributions of representations from both domains projected by the repressor should have a small divergence. Besides, the representations should retain key information that can be used to classify failure samples.

Firstly, all samples are projected to a d -dimensional space by the repressor. The representations of X^s and X_H^t are denoted as $f_r(X^s)$ and $f_r(X_H^t)$, where f_r is the mapping function of the repressor. The distance estimator is then used to measure the distribution divergence of representations from source domain \mathbb{R}^s and minority disk \mathbb{R}_H^t . Here, we introduce Wasserstein distance as the measurement metric because it can measure the divergence between two arbitrary distributions even if they are distant.

Based on Kantorovich Rubinstein theorem, the dual representation of the first Wasserstein distance of two Borel probability measures \mathbb{P} and \mathbb{Q} can be formalized as

$$W_1(\mathbb{P}, \mathbb{Q}) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x \sim \mathbb{P}}[f(x)] - \mathbb{E}_{x \sim \mathbb{Q}}[f(x)] \quad (2)$$

where L -Lipschitz condition is defined as $\|f\|_L = \frac{\sup |f(x) - f(y)|}{\rho(x, y)} \leq L$. Accordingly, the Wasserstein distance of source domain and minority disk $W_1(\mathbb{R}^s, \mathbb{R}_H^t)$ in the latent space can be calculated as:

$$W_1(\mathbb{R}^s, \mathbb{R}_H^t) = \sup_{\|f_d\|_L \leq 1} \mathbb{E}_{x \sim \mathbb{P}_s}[f_d(f_r(x))] - \mathbb{E}_{x \sim \mathbb{P}_{t_H}}[f_d(f_r(x))] \quad (3)$$

where f_d is the function learned by the distance estimator with its parameters θ_d to map representations h to real numbers. Then, we can approximate the empirical Wasserstein distance of representation distribution of source and target domain via maximizing domain critic loss \mathcal{L}_{wd} with respect to θ_d :

$$\mathcal{L}_{wd} = \frac{1}{n_s} \sum_{x^s \in X^s} f_d(f_r(x^s)) - \frac{1}{n_t} \sum_{x^t \in X_H^t} f_d(f_r(x^t)) \quad (4)$$

Note that the f_d should satisfy the 1-Lipschitz condition when calculating the first Wasserstein distance. Therefore, a gradient penalty term $\mathcal{L}_{grad} = \mathbb{E}_{h \sim [\mathbb{R}^s, \mathbb{R}_H^t]}[(\|\nabla_h f_d(h)\|_2 - 1)^2]$ is added to \mathcal{L}_{wd} . And the final objective function of distance estimator (\mathcal{L}_{dist}) can be written as:

$$\max_{\theta_d} \{\mathcal{L}_{wd} + \lambda \mathcal{L}_{grad}\} \quad (5)$$

where the λ is used to balance the \mathcal{L}_{wd} and \mathcal{L}_{grad} .

The failure predictor is used to predict failures for the source disk models. Its inputs are representations from source domain samples and the labels of representations are consistent with their corresponding source domain samples. The objective function of failure predictor (\mathcal{L}_C) can be formalized as:

$$\min_{\theta_c} \frac{1}{N} \sum_{h \in \mathbb{R}^s} -[y_i \cdot \log(f_c(h_i)) + (1 - y_i) \cdot \log(1 - f_c(h_i))] \quad (6)$$

The task of the representor is to ① reduce the representation divergence of the minority disk and source domain disk models, and ② extract the discriminative information from samples in representation learning. Therefore, the objective of representor (\mathcal{L}_R) can be formalized as:

$$\min_{\theta_r} \{\mathcal{L}_C + \gamma \mathcal{L}_{wd}\} \quad (7)$$

where γ is the coefficient to balance between discriminative and transferable feature learning.




Fig. 2. Visualization of the representations from the target domain and the source domain at different confidence thresholds (i.e., conf-level-*)

4.3 Confidence-Based Sample Selection

While introducing more disk models into the source domain helps bring in more failure modes, it also complicates the distribution of the source domain, leading to negative transfer in representation learning. Therefore, it is necessary to filter out irrelevant samples from the source domain to avoid the probable negative transfer.

Specifically, DiskDA achieves this via a confidence-based sample selection process. The confidence of source domain samples is measured based on the similarities of their representations to those from the minority disk. The confidence is given by the domain classifier which is a supervised learning-based binary classifier. The inputs are representations of samples from both domains and the labels are 0/1. The representations from the source domain are labeled as ‘0’ and those from the target domain are labeled as ‘1’. To train a domain classifier, we randomly select a small proportion of representations of samples from the source domain and minority disk. And the objective of the domain classifier can be formulated as:

$$\min_{\theta_D} \frac{1}{N} \sum_i -[y_i \cdot \log(f_c(h_i)) + (1 - y_i) \cdot \log(1 - f_c(h_i))] \quad (8)$$

where $x \in [\mathbb{R}^s, \mathbb{R}_H^t]$, $y \in \{0, 1\}$.

When the parameters converge, the domain classifier is applied to the adjusted source domain as a filter. Since the domain classifier is realized by a deep neural network and the results are given by a sigmoid function in the last layer, the output of the domain classifier is the probability that a representation belongs to the minority disk. So we use this probability to measure the confidence of a sample. DiskDA discards the samples in the source domain whose confidence is lower than the pre-defined threshold and uses the remaining samples in the source domain with higher confidence for further representation learning. By filtering out the ‘low quality’ samples, the distribution divergence between the source domain and the minority disk is reduced.

In Fig. 2, we visualize the distribution of sample representations for the minority disk and the source domain under different confidence thresholds via t-SNE, where the colored points denoted as “conf-level- $*$ ” represent the representations of source domain samples with different confidence values and the black point represent the representations of minority disk samples. t-SNE is short for t-Distributed Stochastic Neighbor Embedding which allows us to project high-dimensional embedding spaces into 2D spaces for visualization while keeping their relative distance. In other words, the points close in the figure have a small distance in the original space. Since the dimension of representations has been compressed, the x and y axes of points have no specific meaning. As seen, the representations of minority disk samples are located closer to that of the source domain samples when a higher confidence threshold is selected, which indicates the rationality of our sample selection.

5 Experiment

In this section, we conduct experiments to evaluate the performance of DiskDA. We first describe the methodology and then show the experimental comparison results among DiskDA and 7 state-of-the-art solutions. Finally, we show the results of sensitivity analysis to explore how the critical hyper-parameters affect the failure prediction performance of DiskDA.

5.1 Methodology

Datasets. The disk models used in our experiments are from two real-world datasets. We select ST4000DM000 (Disk_1), ST6000DX000 (Disk_2), ST3000DM001 (Disk_3), Hitachi HDS5 C4040ALE630 (Disk_4), Hitachi HDS722020ALA330 (Disk_5), Hitachi HDS723030ALA640 (Disk_6), HGST HMS5C4040BLE640 (Disk_7), HGST HMS5C4040ALE640 (Disk_8), HGST HUH728080ALE600 (Disk_9) in Backblaze¹ with a period from 2015-01-01 to 2019-12-31. We select MC1 (SSD_1), MC2 (SSD_2) and MA1 (SSD_3) from Alibaba Cloud [21] with a period from 2019-01-01 to 2019-12-31. All disk models are selected randomly. Each record in both datasets is labeled as healthy or failed on a daily basis.

Attribute Selection. Not all SMRAT attributes are useful for disk failure prediction, we select SMART 1,4,7,12,190,192,193,194,196,197,199 for HDD failure prediction and SMART 1,5,9,12,171,172,174,175,183,190,232,233 for SSD failure prediction via correlation coefficient analysis. Min-max normalization (i.e. $x_{norm} = \frac{x-x_{min}}{x_{max}-x_{min}}$, where x is the raw value of the SMART attribute, x_{max} and x_{min} are the maximum and minimum values of the SMART attribute in the training set) is used to normalize the values of different SMART attributes.

¹ <https://www.backblaze.com/b2/hard-drive-test-data.html>.

Experiment Setup. Regarding records, close to actual failure, will disturb the failure prediction model, a commonly used approach is to label k continuous healthy samples before the actual failure as failure records too [23, 24]. And k is determined via change-point detection and set to be 3 in our experiments. The representation vector length is set to 128 for both representation ability and cost saving (detailed in Sect. 5.2). The coefficient of gradient penalty term λ is set to 10, which is consistent with the setting commonly used in ML models based on Wasserstein distance [1, 14]. The parameter γ used to balance the weight of discriminative and domain invariant representation learning is set to $1e-2$, which is determined through grid search. In each domain invariant representation learning iteration, the distance estimator runs 10 steps then the parameters of the representor and the failure predictor update once. The sample selection process runs one time every 500 representation learning iterations. And the confidence threshold is set to 0.2, which delivers the optimal transfer learning performance (detailed in Sect. 5.2).

Evaluation Metrics. The failure prediction rate (FDR, also called recall), false alarm rate (FAR, also called false positive rate), and F1-Score are adopted as the metrics to measure the disk failure prediction performance. A good disk failure prediction method should reach a high FDR with a low FAR. And the F1-Score is the balance between the FDR and FAR, thus is the most important metric to measure the performance of the prediction model.

Benchmarks. We test three types of benchmarks.

- **Supervised Learning-Based:** We first measure the performance of three supervised learning-based failure prediction methods (i.e. GBRT [10], HDDse [23] and RGF [2]) with only minority disk samples.
- **Semi-supervised Learning-Based:** We explore the performance of the semi-supervised learning model (VAE-LSTM [25]) which models the healthy samples and classifies the failure samples by comparing the reconstruction error with a pre-defined threshold.
- **TL-Based:** We evaluate 3 state-of-the-art TL-based failure prediction models (i.e. TLDFP [24], SSDB [2] and FLBT [15]). Note that TLDFP and SSDB are instance-based TL approaches, and FLBT is a feature-based TL approach. Note that we also test ADA-CBAN [9] in our experiments while its performance is not stable. 3 (i.e., Exp_1, Exp_2, and Exp_4) of the 5 experiments failed due to the large distribution divergence of the source and target domain. In addition, the results of the ADA-CBAN in the last two experiments are worse than those of DiskDA, so we omit to show its results in our comparison. Since all the TL-based methods require a base failure prediction model, a bidirectional gated recurrent unit network (Bi-GRU) is adopted as the base failure prediction model in this paper. Note that the base model can also be replaced by any neural network-based failure prediction model. In fact, we choose Bi-GRU because it shows simplicity, robustness, and accuracy in experiments.

Table 2. The Details about Dataset Used in Experiments.

No.	Training Set	Testing Set
Exp.1	Source: Disk_1, Disk_2 Target: Disk_3 (110, 6)	Disk_3 (1058, 100)
Exp.2	Source: Disk_4, Disk_5 Target: Disk_6 (100, 4)	Disk_6 (918, 40)
Exp.3	Source: Disk_7, Disk_8 Target: Disk_9 (100, 3)	Disk_9 (985, 15)
Exp.4	Source: Disk_8, Disk_9 Target: Disk_4 (300, 4)	Disk_4 (2360, 38)
Exp.5	Source: SSD_1, SSD_2 Target: SSD_3 (399, 13)	SSD_3 (39563, 1357)

Considering that the three comparative TL-based methods are all single source-based TL models, we traverse all source domain disk models and select the one with the best performance as the source domain.

The detailed setup of training and testing datasets is shown in Table 2. Note that Disk.a (x, y) denotes that the number of healthy disks is x and the number of the failed disk is y. Generally, the number of healthy samples in a real dataset is much larger than that of failure samples. In order to avoid model bias caused by data imbalance, we only use the randomly selected healthy samples with the same number of failure samples. Cross-validation is done for each method, and we average their performance as the final result.

Table 3. Performance Comparison of Various Disk Failure Prediction Models

Method	Exp.1			Exp.2			Exp.3			Exp.4			Exp.5		
	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1
RGF	0.382	0.021	0.536	0.698	0.056	0.725	0.621	0.017	0.726	0.618	0.06	0.557	0.747	0.034	0.793
GBRT	0.605	0.106	0.666	0.674	0.035	0.744	0.5	0.014	0.637	0.532	0.017	0.624	0.684	0.013	0.801
HDDse	0.634	0.057	0.702	0.784	0.009	0.891	0.715	0.013	0.757	0.573	0.023	0.649	0.624	0.003	0.799
VAE-LSTM	0.756	0.103	0.807	0.821	0.095	0.886	0.796	0.031	0.848	0.697	0.046	0.723	0.785	0.067	0.815
Bi-GRU	0.491	0.098	0.582	0.684	0.011	0.792	0.68	0.021	0.761	0.355	0.04	0.405	0.51	0.004	0.672
DiskDA	0.917	0.027	0.976	0.991	0.002	0.998	0.913	0.003	0.974	0.907	0.001	0.969	0.973	0.043	0.993

5.2 Experimental Results

Performance Comparison with Supervised and Semi-supervised Learning Methods. We first show the performance comparison results of




Fig. 3. Visualization of representation from samples of minority disk and source domain disk models

DiskDA with supervised learning methods (RGF, GBRT, and HDDse) and semi-supervised learning methods (VAE-LSTM). The training set used is the dataset indicated as “Target” in Table 2.

It can be seen from Table 3 that

- The supervised learning-based methods perform poorly due to insufficient training data. Although HDDse adopts a metric learning method that actually increases the size of the training set by taking pairs of samples as input, its prediction performance is only slightly better than the other two.
- Compared with the above supervised-learning methods, the FDR value of VAE-LSTM is significantly improved. However, because the threshold-based method adopted by VAE-LSTM cannot well classify the failure samples from the healthy, the FAR is also higher.
- DiskDA has the best failure prediction performance, with an F1-Score 20.86% higher than VAE-LSTM on average. In addition, we also find that the F1-Score of DiskDA is 61.75% higher than that of the base model, which also proves the necessity and effectiveness of domain adaption. DiskDA can reach the best performance because ① it can extract failure prediction knowledge from large amounts of source domain disk samples rather than the limited samples provided by the minority disk; ② it adopts a supervised learning-based approach that can automatically discriminate the healthy and failure samples without setting classification threshold manually.

Performance Comparison with TL-Based Methods. Next, we show the performance comparison results of DiskDA with TL-based methods (SSDB, TLDFP, and FLBT). Since all these TL-based methods can only work if there are failure samples in the target domain, DiskDA also uses failure samples in the

target domain as other methods for fair comparison (see Table 2 for details of the training set). However, the DiskDA can work without failure samples, which is superior to other solutions. To highlight that, we further implement DiskDA in a more restrictive case where no failure sample of minority disks is provided and the results are indicated as “DiskDA*”. To facilitate analysis, we also test the performance of the base model trained in the source domain on the testing set as the baseline (indicated as “Src”). From Table 4 we can see that

- The F1-Score of Src is only 0.592 on average. The root cause is the distribution shift of SMART attributes among disks, so simply reusing the prediction model trained upon other disks will fail in practice.
- Among the three TL-based methods, the instance-based TL approach TLDFP has the best performance. By continuously enhancing the weight of misclassified samples, the ability of the failure predictor can be improved to a certain extent. Another instance-based TL approach, SSDS, uses a sample selection strategy based on similarity to adjust the source domain samples. However, its performance is even worse than the baseline in some cases. The performance of FLBT is worse than that of TLDFP. All the TL-based approaches can just reach sub-optimal performance as they drop useful samples and attributes and can only utilize partial information from the source domain.

Table 4. Performance Evaluation of Transfer Learning Based Failure Prediction Models

Method	Exp.1			Exp.2			Exp.3			Exp.4			Exp.5		
	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1	FDR	FAR	F1
Src	0.486	0.09	0.584	0.979	0.605	0.618	0.916	0.762	0.537	0.97	0.919	0.505	0.977	0.472	0.715
SSDB	0.34	0.467	0.302	0.911	0.384	0.801	0.991	0.84	0.707	0.95	0.68	0.685	0.71	0.672	0.597
TLDFP	0.785	0.094	0.836	0.751	0.001	0.858	0.811	0.422	0.773	0.648	0.001	0.786	0.804	0.117	0.843
FLBT	0.685	0.194	0.636	0.747	0.131	0.834	0.611	0.134	0.683	0.588	0.112	0.716	0.722	0.093	0.837
DiskDA*	0.84	0.048	0.889	0.987	0.002	0.994	0.852	0.005	0.92	0.825	0.002	0.904	0.962	0.092	0.947
DiskDA	0.917	0.027	0.976	0.991	0.002	0.998	0.913	0.003	0.974	0.907	0.001	0.969	0.973	0.043	0.993

- DiskDA performs the best, with the F1-Score 20.02% higher than the best competitor. This is because ① DiskDA can extract failure prediction knowledge from multiple source domain disk models while existing instance-based TL approaches can only benefit from a single source domain (just one disk model in the source domain) to prevent negative transfer. ② DiskDA tries to fully utilize the source domain disk samples by fusing the distribution of source and target domain samples in the latent space, rather than directly dropping samples or attributes as existing TL-based approaches. We visualize the fusion of representations from the source and target domain via t-SNE in Fig. 3, where the red points represent the representations from the source domain and the green ones represent the representations from the target domain. As seen, the green and red points are fused constantly as the




Fig. 4. Performance of introducing more disk models to source domain

increment of iterations, which suggests that the domain invariant representation learning process can effectively fuse the representations from the source and target domain and thus the failure prediction knowledge extracted from the source domain can adapt to the target domain(minority disk).

- In the absence of failure samples from the minority disk, the F1-Score of DiskDA can still reach a satisfactory 0.93, 66.75% higher than that of the baseline. The results prove the effectiveness of our theorem that DiskDA can still adapt the failure prediction knowledge extracted from source domain disk models to minority disks with no failure sample. And this shows that DiskDA has much higher adaptivity compared to existing TL-based approaches. We further explore whether DiskDA can continuously benefit from the increment of source domain failure modes. And we investigate this problem by measuring the performance of DiskDA as the increase of source domain disk models because more disk models potentially contain more failure modes. As shown in Fig. 4, the performance of DiskDA can be continuously improved by adding more disk models to the source domain. Thanks to the sample selection process, DiskDA can timely filter out the source domain samples deteriorating the domain invariant representation learning process and effectively transfer failure prediction knowledge from the source domain to the minority disk. The results motivate us to add more disk models to the source domain to reach high performance without worrying about the negative transfer problem.

5.3 Sensitivity Study

Impact of Hidden Size. In domain invariant representation learning, the representor projects samples from both domains to fixed-length vectors as their representations. And we evaluate how the size of length affects the performance of DiskDA. As seen in Fig. 5.a, the performance of DiskDA is improved as the increase of hidden size and then steady until the hidden size reaches 128. The results indicate that a small size will limit the representation ability and seriously affect the performance of DiskDA. And we set the hidden size as 128 to reach a balance between the performance and computation cost.




Fig. 5. Sensitivity Study

Impact of Confidence-Threshold. The confidence threshold determines which samples to select in representation learning afterward. And we explore how different confidence thresholds affect the performance of DiskDA. In the experiment, we experiment with different confidence threshold values and summarize the results as 3 representative curves corresponding to thresholds of 0.4 (red curve), 0.2 (green curve), 0 (blue curve), shown in Fig. 5.b. As seen, a low (i.e., ranging from 0 to 0.1) or high threshold (i.e., ranging from 0.4 to 1) can both deteriorate the performance of DiskDA since the negative transfer caused by irrelevant samples remaining in the source domain or the loss of relevant samples filtered out in confidence-based sample selection process. And we set the threshold as 0.2 to achieve a balance between the performance of domain adaption and the loss of source domain samples.

6 Conclusion

In this work, we investigate the problem of minority disk failure prediction in data centers. Based on the fact that the failure modes are common for different disk models, our basic idea is to utilize full of the failure prediction knowledge learned from other disk models to the minority disk. We model this as an unsupervised domain adaption problem and analyzed the generalization error bound of the prediction model in the target domain (minority disk) theoretically. Guided by the generalization error bound, we design a framework which can effectively optimize the error bound by elaborately combining the domain invariant representation and confidence-based sample selection processes. Our experiments on real-world datasets show the effectiveness of our approach. Moreover, our approach can still reach a satisfying F1-score of 0.93 on average for minority disks even with no failure samples, which suggests that our approach can fit for more broad cases compared to existing approaches.

7 Appendix

Proof. The discrepancy between the source and target domain is measured using the Wasserstein distance in DiskDA. Specifically, the p -th Wasserstein distance between two Borel probability measures \mathbb{P} and \mathbb{Q} is defined as:

$$W_p(\mathbb{P}, \mathbb{Q}) = \left(\inf_{\mu \in \Gamma(\mathbb{P}, \mathbb{Q})} \int \rho(x, y)^p d\mu(x, y) \right)^{1/p} \quad (9)$$

where the $\Gamma(\mathbb{P}, \mathbb{Q})$ is the set of all joint distributions $\mu(x, y)$ whose marginal distribution are \mathbb{P} and \mathbb{Q} . The $\mu(x, y)$ can be viewed as a policy for transporting a unit quantity of material from x to y and the $\rho(x, y)$ is the corresponding cost. And the Wasserstein distance between \mathbb{P} and \mathbb{Q} represents the minimum expected transport cost. As Wasserstein distance satisfies the triangle inequality, the following equation holds

$$W_p(\mathbb{P}_s, \mathbb{P}_t) \leq W_p(\mathbb{P}_s, \mathbb{P}_{t_H}) + W_p(\mathbb{P}_{t_H}, \mathbb{P}_t) \quad (10)$$

Shen et al. [14] prove the generalization error bound of a classification function h in the target domain for unsupervised domain adaption based on Wasserstein distance as

$$\epsilon_t(h) \leq \epsilon_s(h) + 2KW_1(\mathbb{P}_s, \mathbb{P}_t) + \lambda \quad (11)$$

where the K means that all hypotheses h are K -Lipschitz continuous, λ is the combined error of the optimal hypothesis h^* which minimizes the combined error $\epsilon_s(h) + \epsilon_t(h)$, \mathbb{P}_s and \mathbb{P}_t are distributions of source and target domain, respectively. Let C denote $2KW_1(\mathbb{P}_{t_H}, \mathbb{P}_t)$. By substituting inequality (11) for (10), Theorem 3.1 is derived.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International Conference on Machine Learning, pp. 214–223. PMLR (2017)
2. Botezatu, M.M., Giurgiu, I., Bogojeska, J., Wiesmann, D.: Predicting disk replacement towards reliable data centers. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 39–48 (2016)
3. Chakrabortii, C., Litz, H.: Improving the accuracy, adaptability, and interpretability of SSD failure prediction models. In: Proceedings of the 11th ACM Symposium on Cloud Computing, pp. 120–133 (2020)
4. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005), vol. 1, pp. 539–546. IEEE (2005)
5. Ghemawat, S., Gobiuff, H., Leung, S.T.: The google file system. In: Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles, pp. 29–43 (2003)

6. Jiang, T., Zeng, J., Zhou, K., Huang, P., Yang, T.: Lifelong disk failure prediction via gan-based anomaly detection. In: 2019 IEEE 37th International Conference on Computer Design (ICCD), pp. 199–207. IEEE (2019)
7. Jiang, W., Hu, C., Zhou, Y., Kanevsky, A.: Are disks the dominant contributor for storage failures? A comprehensive study of storage subsystem failure characteristics. *ACM Trans. Storage (TOS)* **4**(3), 1–25 (2008)
8. Johnson, R., Zhang, T.: Learning nonlinear functions using regularized greedy forest. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(5), 942–954 (2013)
9. Lan, X., et al.: Adversarial domain adaptation with correlation-based association networks for longitudinal disk fault prediction. In: 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE (2021)
10. Li, J., et al.: Hard drive failure prediction using classification and regression trees. In: 2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, pp. 383–394. IEEE (2014)
11. Lu, S., Luo, B., Patel, T., Yao, Y., Tiwari, D., Shi, W.: Making disk failure predictions smarter! In: FAST, pp. 151–167 (2020)
12. Mikolov, T., Kombrink, S., Burget, L., Černocký, J., Khudanpur, S.: Extensions of recurrent neural network language model. In: 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 5528–5531. IEEE (2011)
13. Schroeder, B., Gibson, G.A.: Understanding disk failure rates: what does an MTTF of 1,000,000 hours mean to you? *ACM Trans. Storage (TOS)* **3**(3), 8-es (2007)
14. Shen, J., Qu, Y., Zhang, W., Yu, Y.: Wasserstein distance guided representation learning for domain adaptation. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018)
15. Sun, X., et al.: System-level hardware failure prediction using deep learning. In: Proceedings of the 56th Annual Design Automation Conference 2019, pp. 1–6 (2019)
16. Vishwanath, K.V., Nagappan, N.: Characterizing cloud computing hardware reliability. In: Proceedings of the 1st ACM Symposium on Cloud Computing, pp. 193–204 (2010)
17. Wang, Y., Miao, Q., Ma, E.W., Tsui, K.L., Pecht, M.G.: Online anomaly detection for hard disk drives based on mahalanobis distance. *IEEE Trans. Reliab.* **62**(1), 136–145 (2013)
18. Wilson, G., Cook, D.J.: A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.* **11**(5), 1–46 (2020)
19. Xie, Y., Feng, D., Wang, F., Zhang, X., Han, J., Tang, X.: OME: an optimized modeling engine for disk failure prediction in heterogeneous datacenter. In: 2018 IEEE 36th International Conference on Computer Design (ICCD), pp. 561–564. IEEE (2018)
20. Xu, C., Wang, G., Liu, X., Guo, D., Liu, T.Y.: Health status assessment and failure prediction for hard drives with recurrent neural networks. *IEEE Trans. Comput.* **65**(11), 3502–3508 (2016)
21. Xu, F., Han, S., Lee, P.P., Liu, Y., He, C., Liu, J.: General feature selection for failure prediction in large-scale SSD deployment. In: 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pp. 263–270. IEEE (2021)
22. Yang, W., Hu, D., Liu, Y., Wang, S., Jiang, T.: Hard drive failure prediction using big data. In: 2015 IEEE 34th Symposium on Reliable Distributed Systems Workshop (SRDSW), pp. 13–18. IEEE (2015)

23. Zhang, J., Huang, P., Zhou, K., Xie, M., Schelter, S.: HDDSE: enabling high-dimensional disk state embedding for generic failure detection system of heterogeneous disks in large data centers. In: Proceedings of the 2020 USENIX Conference on Usenix Annual Technical Conference, pp. 111–126 (2020)
24. Zhang, J., et al.: Minority disk failure prediction based on transfer learning in large data centers of heterogeneous disk systems. *IEEE Trans. Parallel Distrib. Syst.* **31**(9), 2155–2169 (2020)
25. Zhou, H., et al.: A proactive failure tolerant mechanism for SSDS storage systems based on unsupervised learning. In: 2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS), pp. 1–10. IEEE (2021)