USENIX ATC '20
2020 USENIX Annual Technical Conference
JULY 15–17, 2020

HDDse: Enabling High-Dimensional Disk State Embedding for Generic Failure Detection System of Heterogeneous Disks in Large Data Centers

Ji Zhang[1,4], Ping Huang[1,2], Ke Zhou[1], Ming Xie[3], Sebastian Schelter[4]

[1]Huazhong University of Science and Technology, [2]Temple University

[3]Tencent Inc., [4]University of Amsterdam

iDSM@HUST
INTELLIGENT DATA STORAGE & MANAGEMENT LABORATORY

AI for System

jizhang@hust.edu.cn

http://www.jizhang.pro

Ji Zhang

Chief Engineer, Huawei

Postdoc, University of Amsterdam

Disk Failure Prediction (ATC, ICPP, TPDS)

Sector Error Prediction (DAC, FAST)

Database Tuning (SIGMOD, VLDB, VLDBJ, NEDB)

Disk Failure

Data Protection Scheme

eg., replication and erasure codes

Disk Failure Prediction

eg., Machine Learning-based Disk Failure Prediction

S.M.A.R.T Technology

Self-**M**onitoring **A**nalysis and **R**eporting **T**echnology

(ID, Normalized, Raw, Threshold, Worst)

SMART technology contains up to 30 attributes, reporting various disk operating conditions.

Threshold Method

Failure detection rate (FDR) of 3%-10% with 0.1% false alarm rate (FAR)

Six classes of approaches

- Threshold-Based approaches (TB)
- Distance-based Anomaly Detection approaches (DAD)
- Shallow Machine Learning-based approaches (SML)
- Deep Neural Network-based approaches (DNN)
- One-Class Classification -based approaches (OCC)
- Transfer Learning-based approaches (TL)

Limitations

	TB	DAD	SML	DNN	OCC	TL
Applicability	✓	✓	✗	✗	✗	✗
Adaptability	✓	✓	✗	✗	✗	✓(*)
Imbalance Datasets	✓	✗	✗	✗	✓	✗
Minority Disk	✓	✗	✗	✗	✗	✓(*)
Performance	FDR:3%-10%	FDR:56%-70%	FDR:75%-96%	FDR:87%-98%	FDR:70%-92%	FDR:80%-97%
	FAR:0.1%-2%	FAR:0%-0.8%	FAR:0.8%-4%	FAR:0.6%-1.9%	FAR:0%-10%	FAR:0.5%-6%

(*) refers to certain conditions that are required, e.g., finding a suitable source domain (i.e., another disk model) for knowledge transfer.

Limitation

Applicability and Adaptability

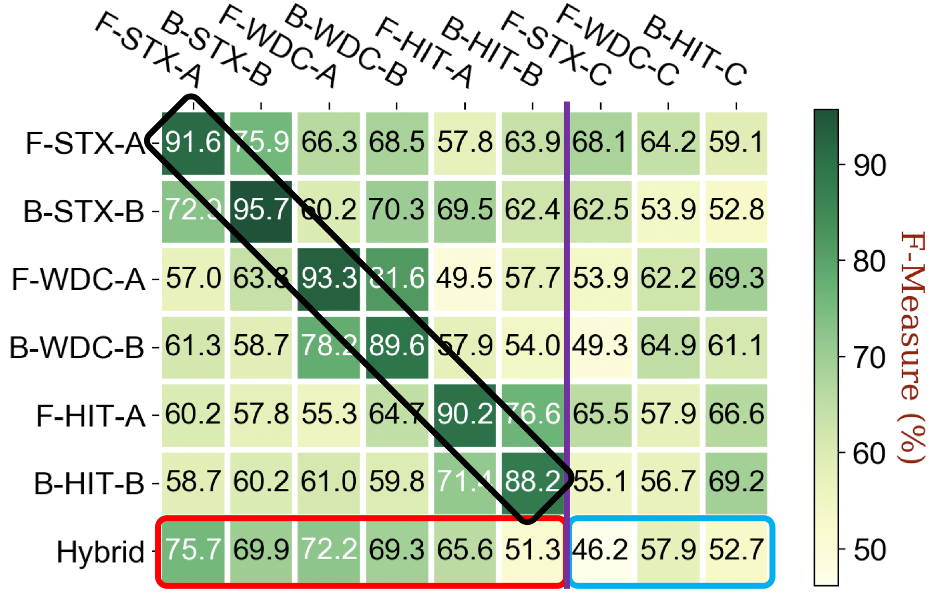

SML

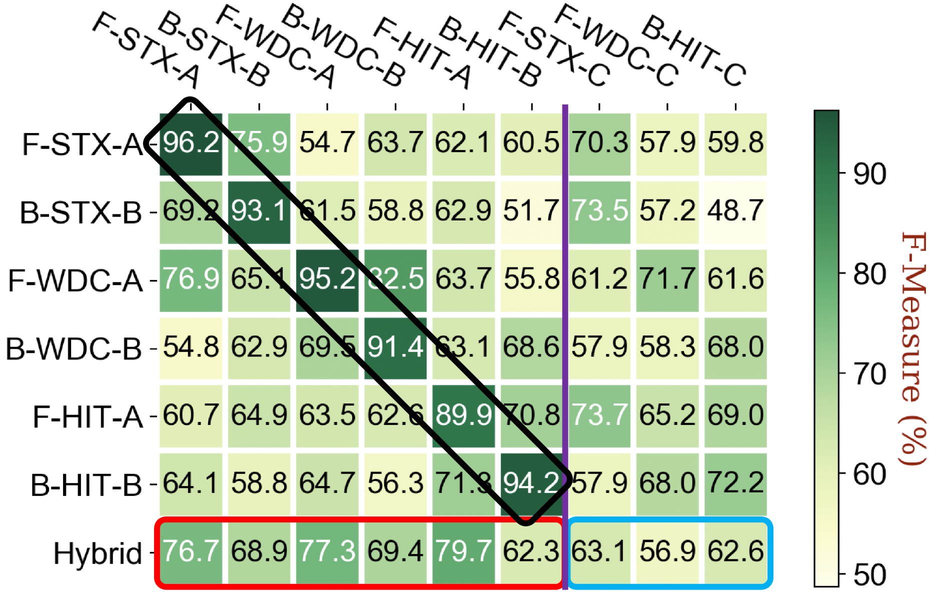

DNN

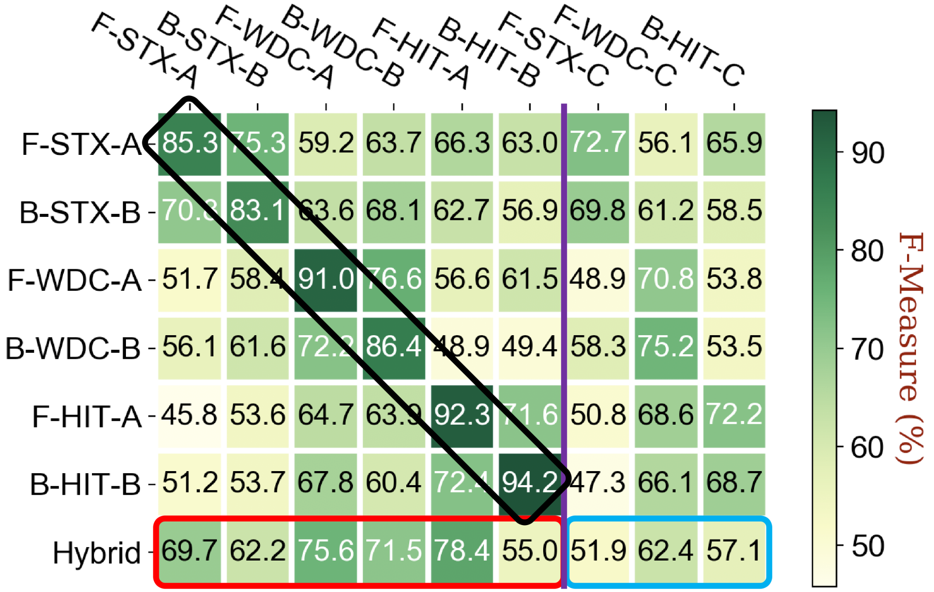

OCC

Limitation

Applicability and Adaptability

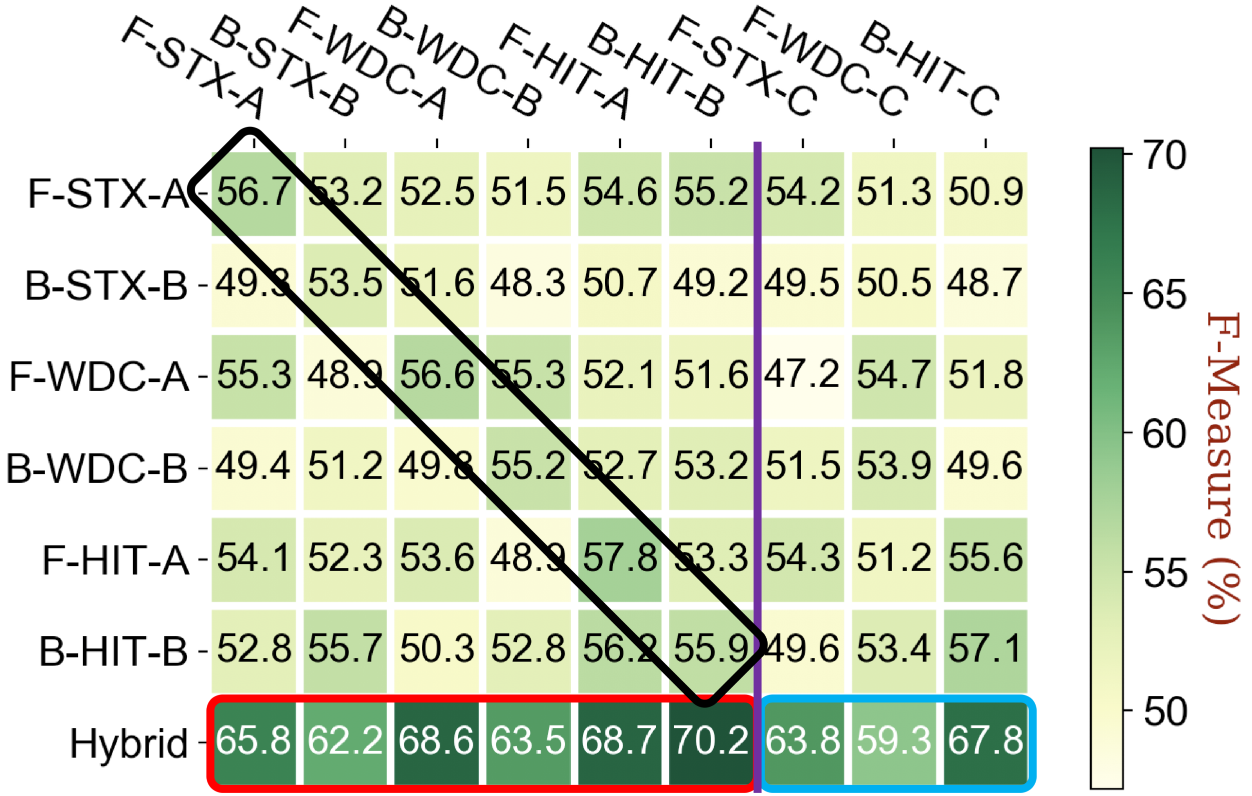

DAD

Limitation

Minority Disk

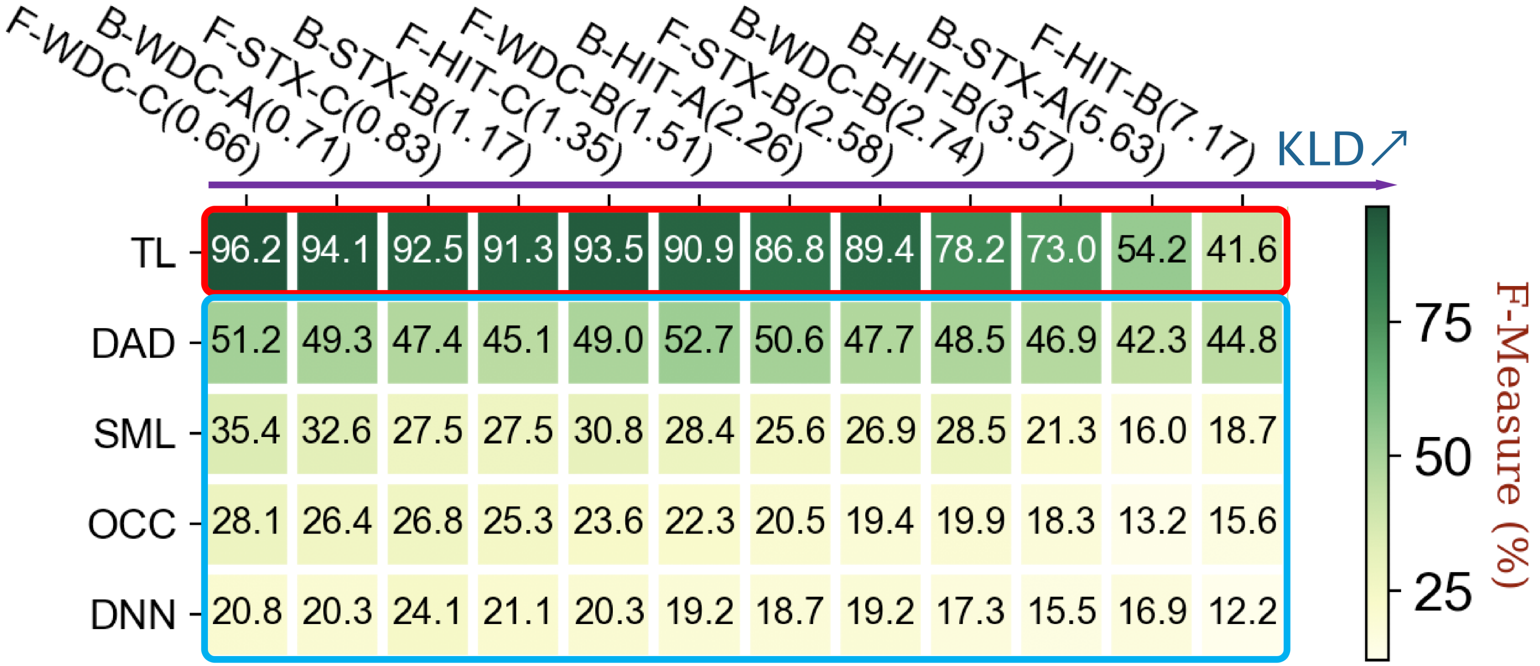

Data Center	KLD(0 ~ 1)	KLD(1 ~ 2)	KLD(2 ~ 3)	KLD(>4)
Tencent	35%	25%	23%	17%
Backblaze	32%	18%	31%	19%

Motivation

- **Why the DAD approaches have good applicability and high adaptability while DNN does not?**
  - A commonality and not sensitive to the disk models.
- **Why the overall detective performance of the DAD method is not as good as other approaches?**
  - Transformation and computation in low-dimensional space.
- **Why the DNN approach achieves the best performance among other candidates**
  - Good expression and fitting ability.

Overview of *HDDse*

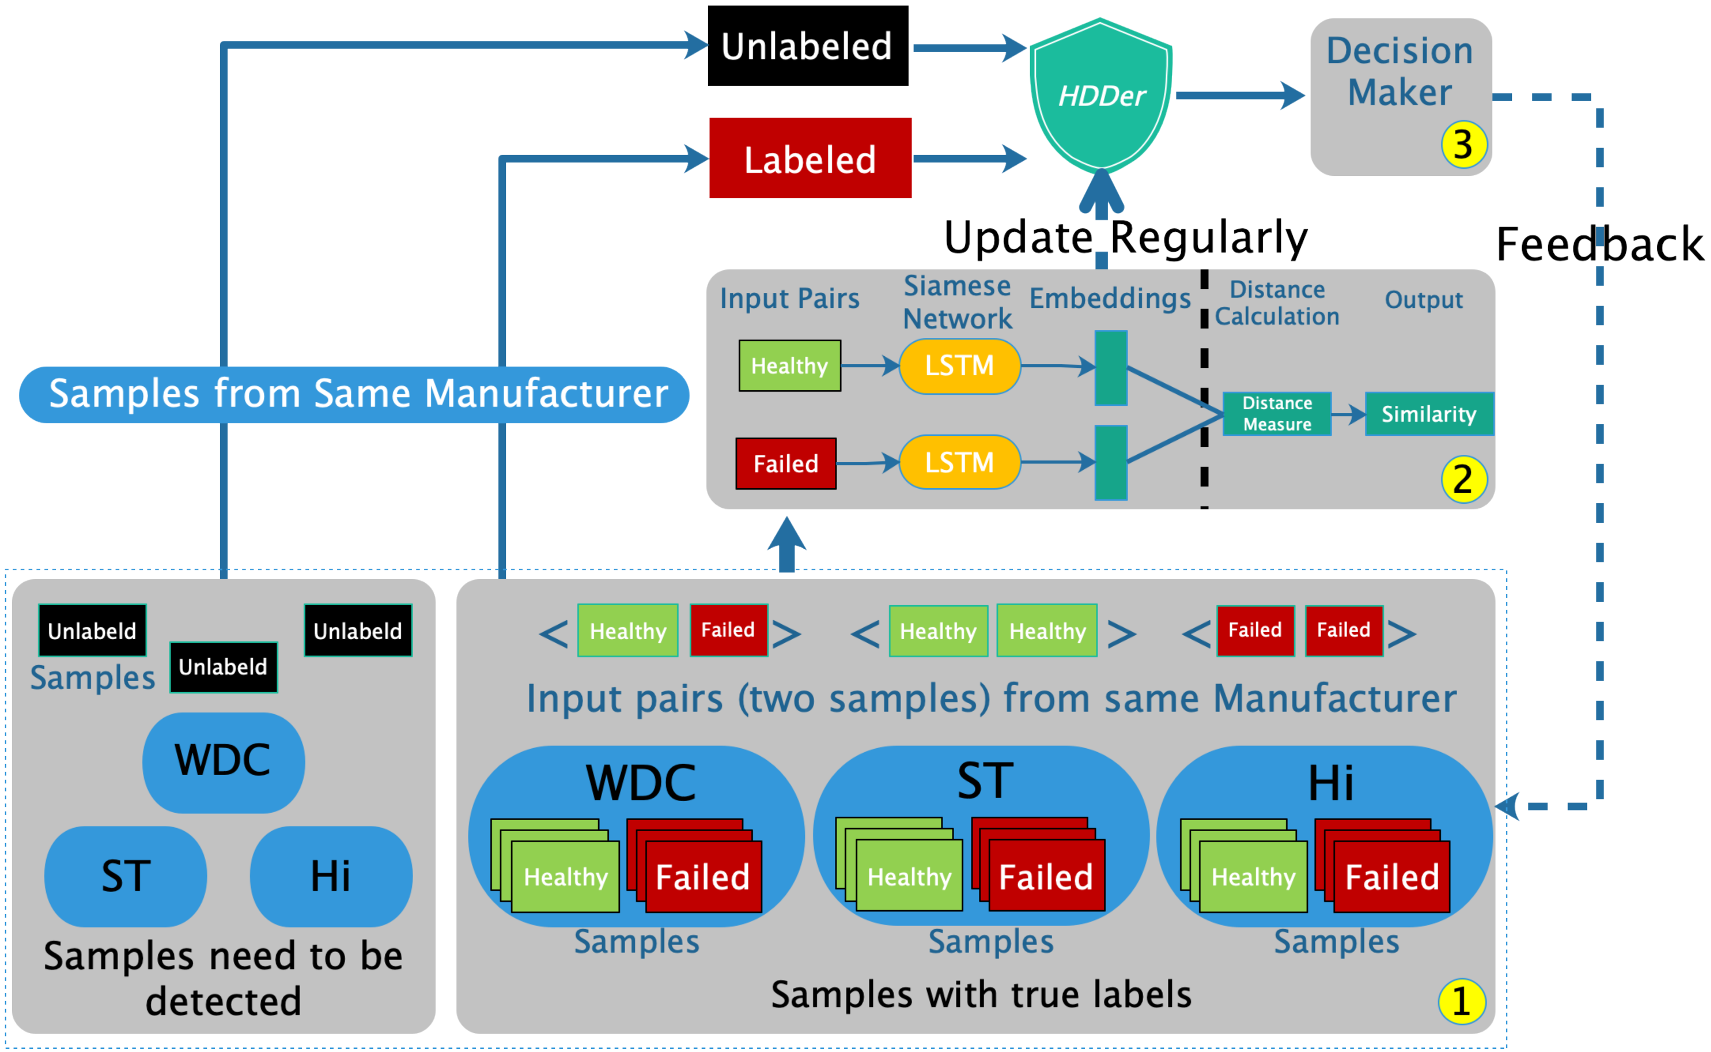

USENIX
ATC '20
2020 USENIX Annual Technical Conference
JULY 15–17, 2020

LSTM-based Siamese Network in ***HDDse***

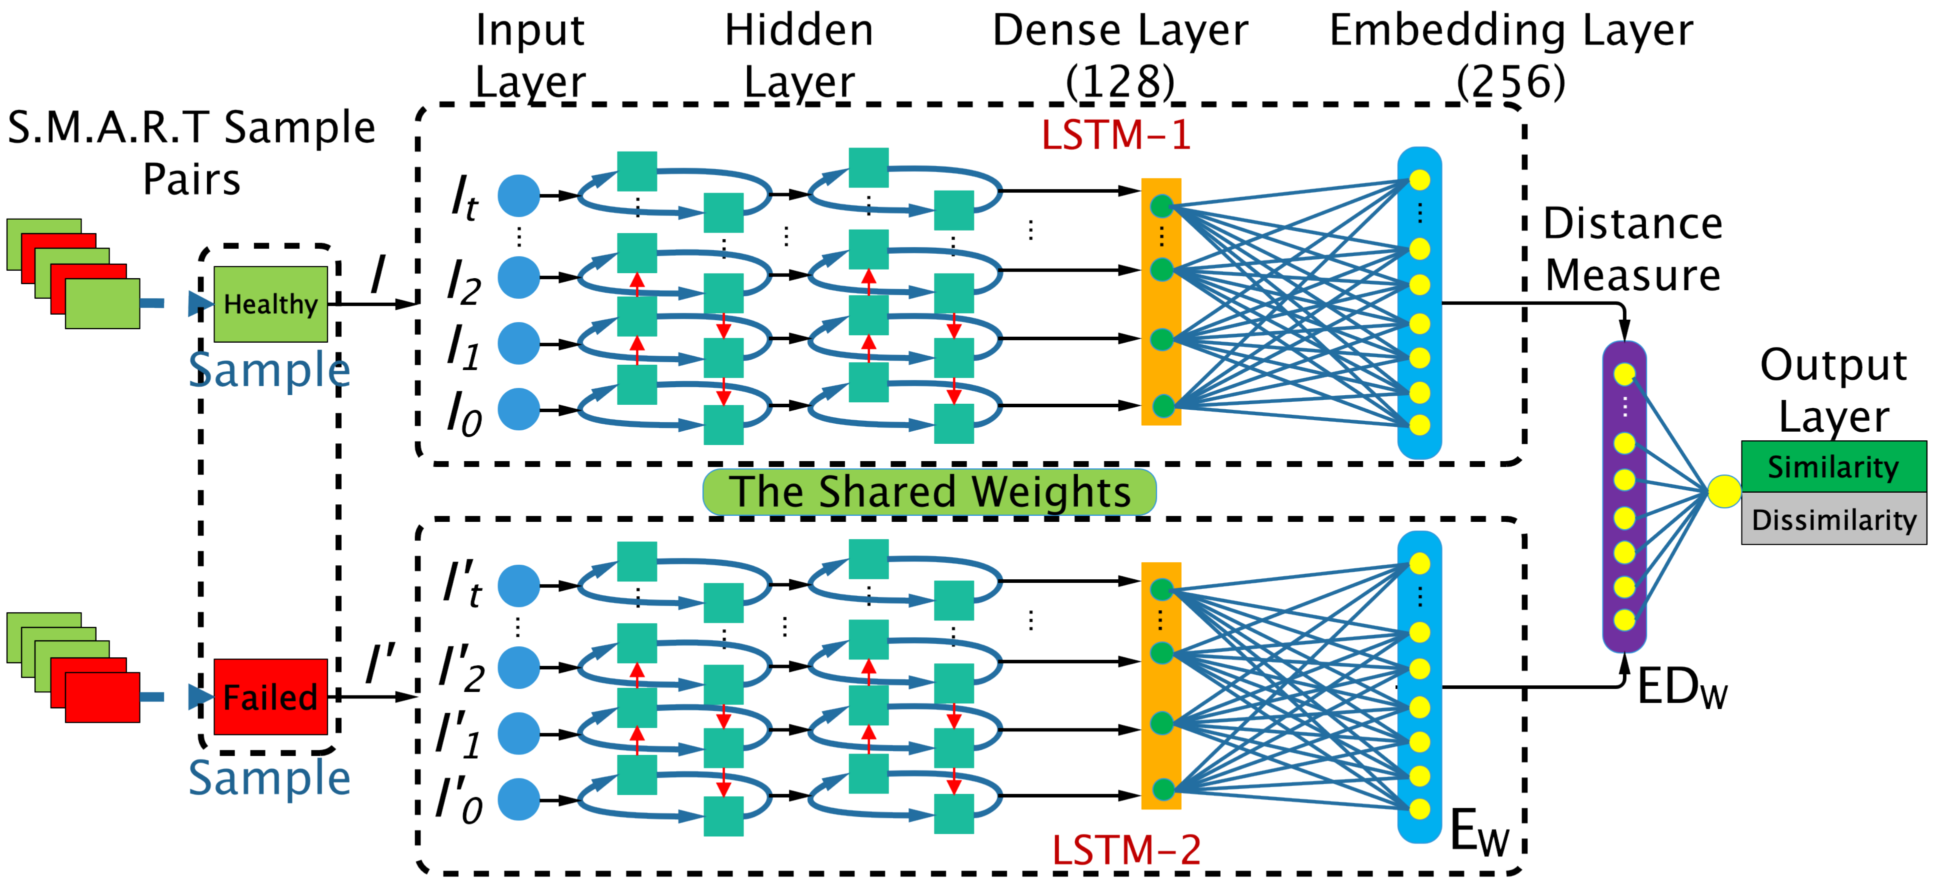

The relationship between instances, samples and the input pairs in ***HDDse***

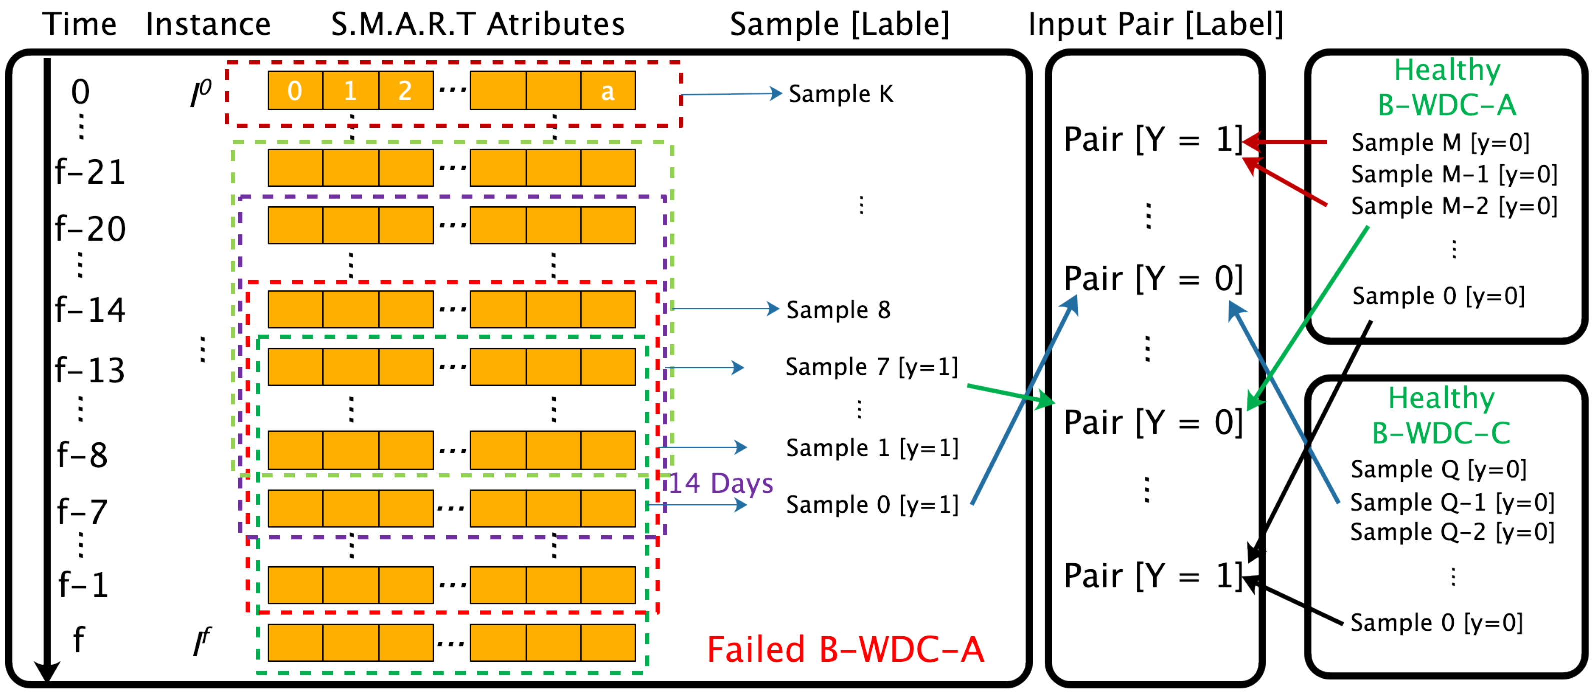

Benefits

Imbalance degree (*IDe*).

For an imbalanced dataset containing a minority class sample with size A and the *IDe* is α , the majority class sample size is α A.

- ## Better with imbalanced datasets

The new imbalance degre $IDe' = \frac{n_1}{n_0} = \frac{\alpha}{2} - \frac{1+\alpha-A}{2A\alpha}$ since $A, \alpha > 1$, which effectively alleviates the original data imbalance by a factor of two.

- ## Better with minority disk models

The number of training pairs with the minority disk models in existing methods is $P = A(1+\alpha)$

In our method the number of training pairs is $\frac{P!}{2!(P-2)!} = \frac{P(P-1)}{2}$

Decision Maker in *HDDse*

A voting-based sliding window (VSW)

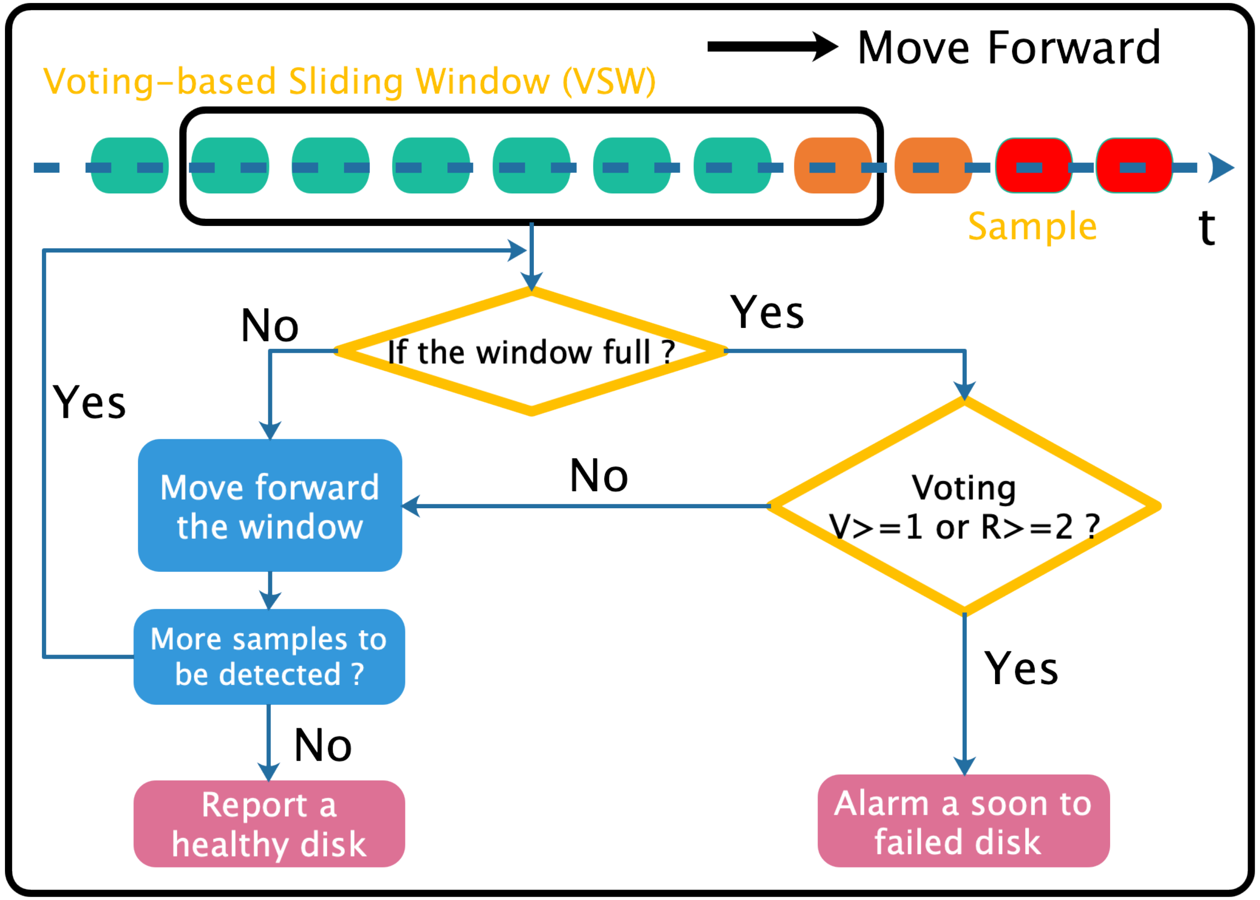

Define a length-W time sliding window and move it forward everyday.

Experimental Evaluation

Datasets:

- From Backblaze, which spans a period of 58 months consisting of 146,203 healthy disks and 8,256 failed disks.
- Tencent and spans 29 months consisting of 70,192 healthy disks and 2,971 failed disks.

Evaluation Metrics:

- **TPR**. True Positive Rate (also called recall) is the proportion of failed disks that are correctly predicted.
- **FPR**. False Positive Rate (also called false alarm rate) is the proportion of healthy disks that are falsely predicted as failed.
- **AUC**. Area under the receiver operating characteristic curve value under the ROC curve (receiver operating characteristic) to evaluate the binary classification performance of our detection model in imbalanced datasets.
- **F-Measure**. A balance between the two metrics TPR and Prediction Precision.
- ***C-MTTDL***. Cost-based MTTDL.

Cost-based Mean Time To Data Loss (C-MTTDL)

$$MTTDL \approx \frac{MTTF}{1-\frac{k\mu}{\mu+\gamma}}$$

approximate the mean time to data loss with failure detection model

Neglect the cost of misclassification by the approach！

An end-to-end economic analysis metric called **C-MTTDL**

$$C-MTTDL = \frac{MTTDL}{Cost} \approx \frac{MTTF}{(1-\frac{k\mu}{\mu+\gamma})(C_aFP+C_bFN)}$$

The t-SNE of the S.M.A.R.T data before and after embedding using ***HDDse***

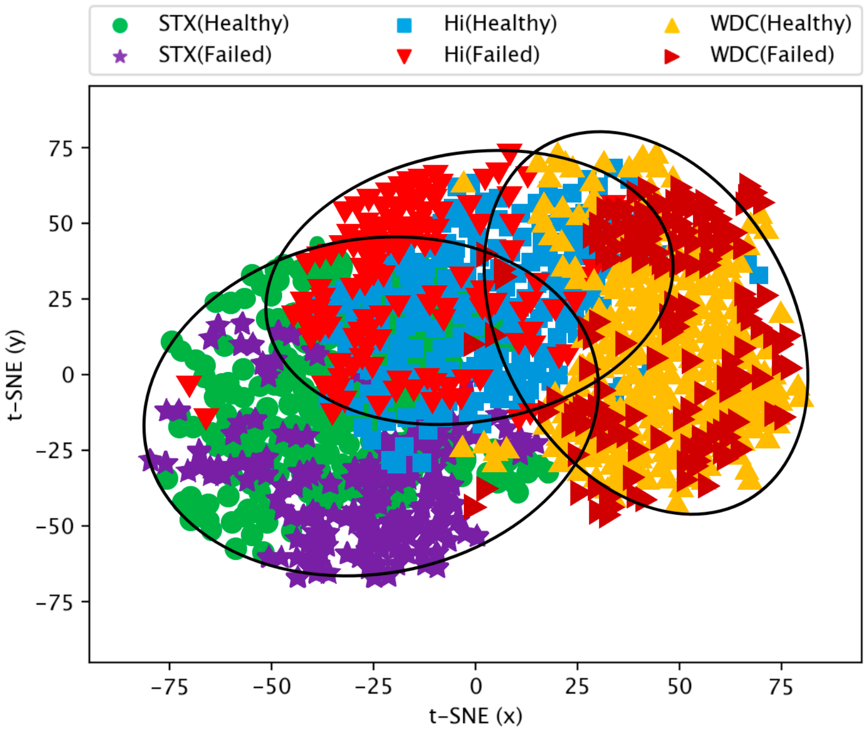

Before

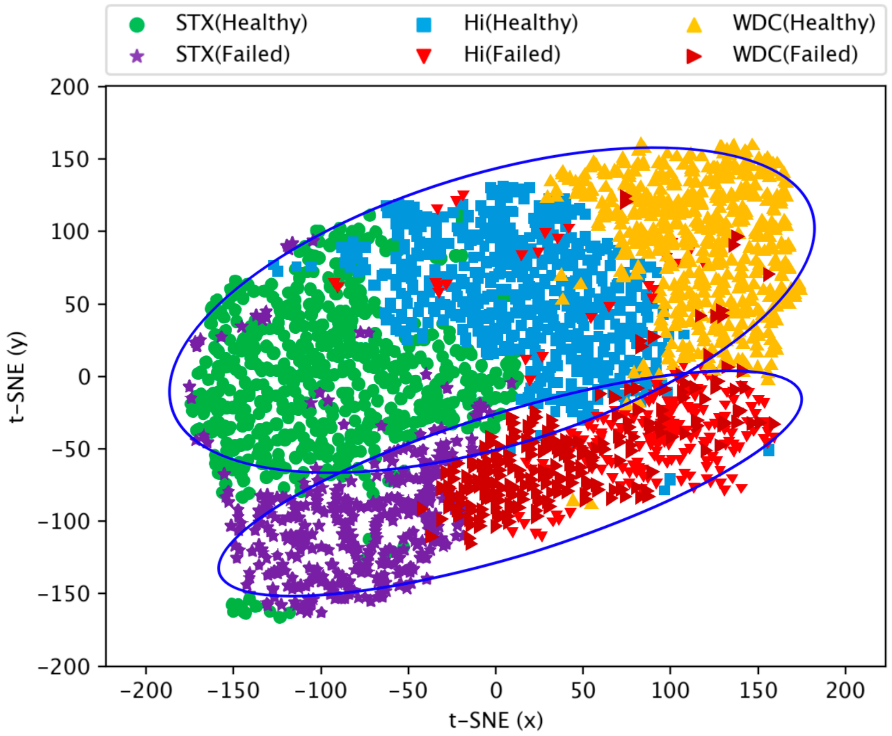

After

HHDse only Trained on Minority Disk Datasets

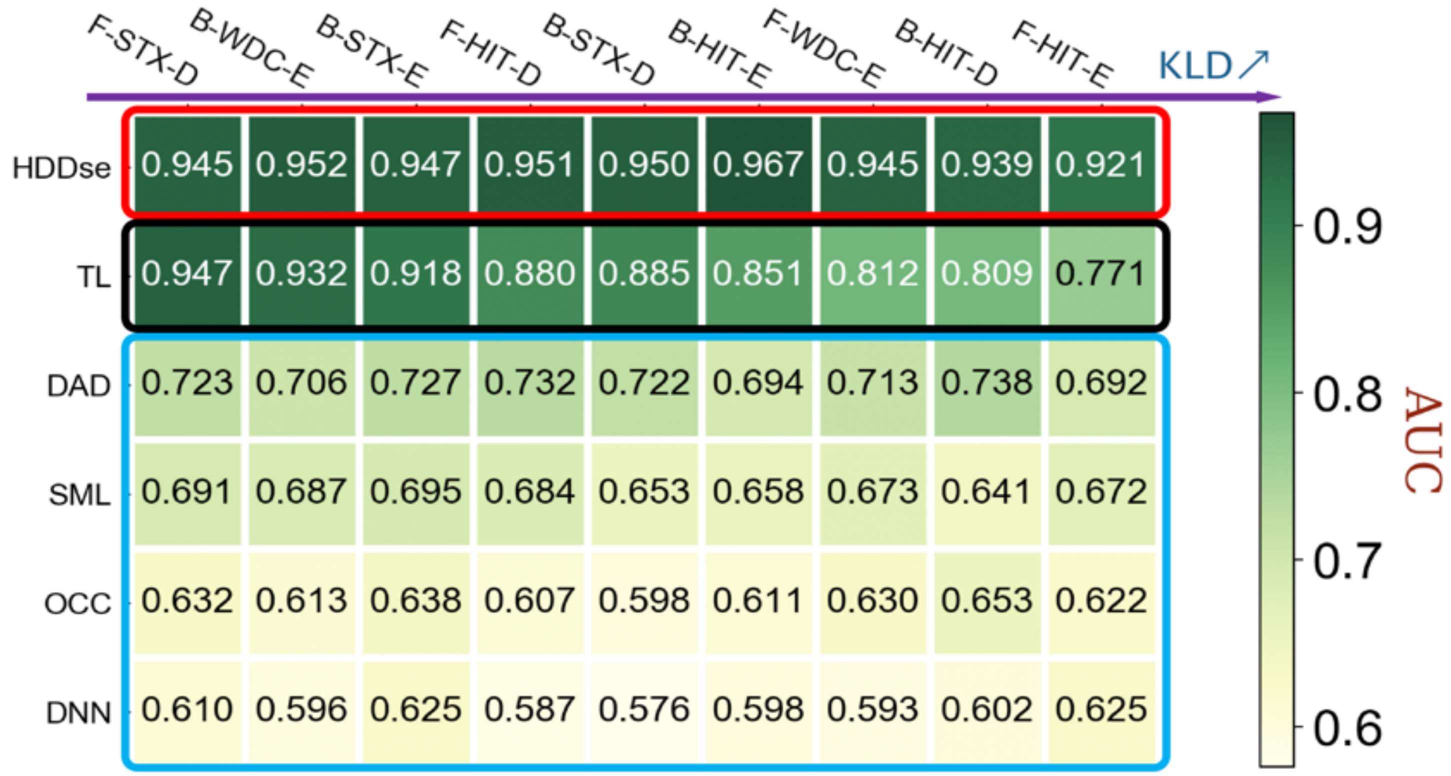

USENIX
ATC '20
2020 USENIX Annual Technical Conference
JULY 15–17, 2020

The Applicability of ***HHDse***

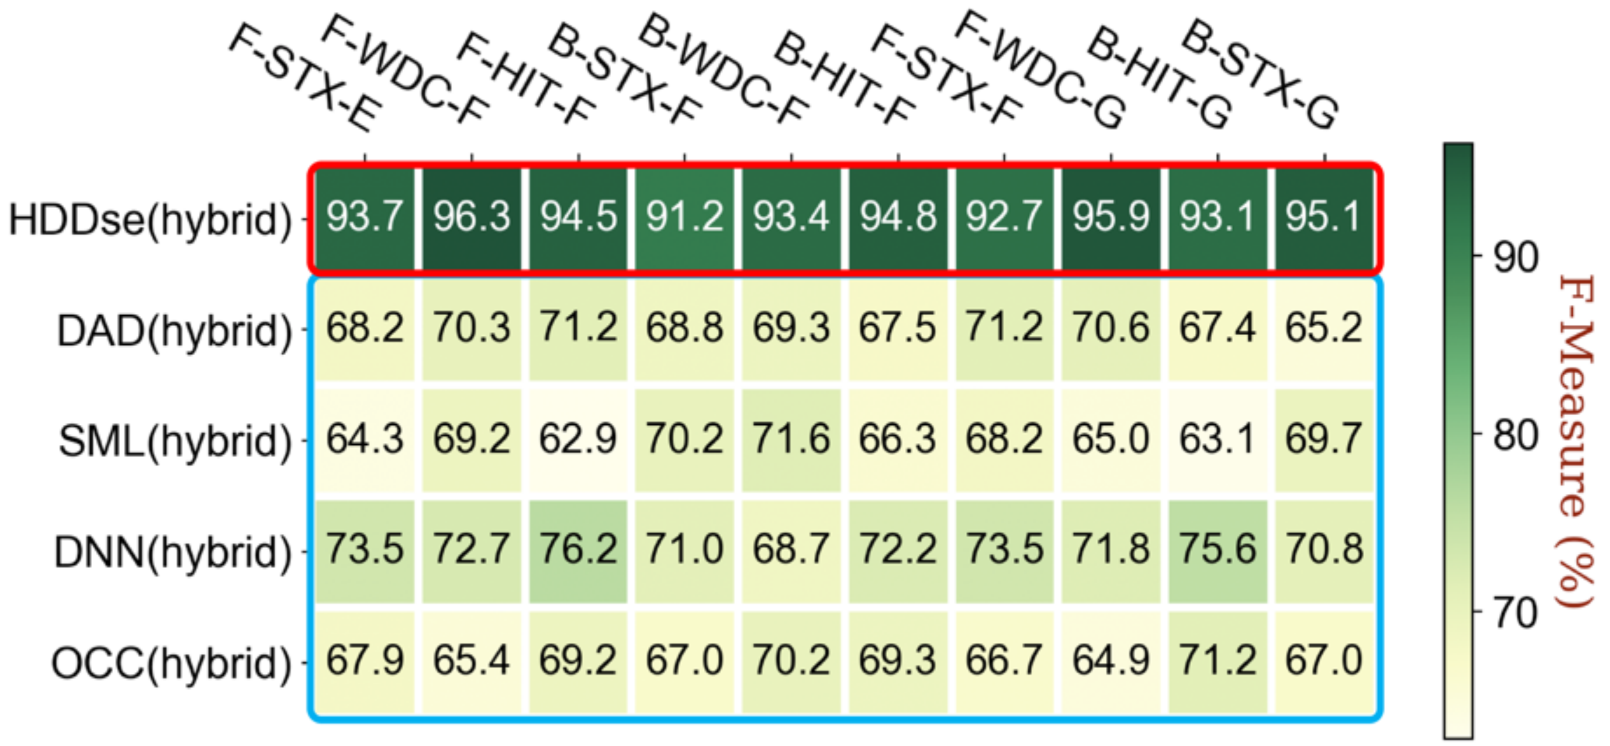

The Adaptability of ***HHDse***

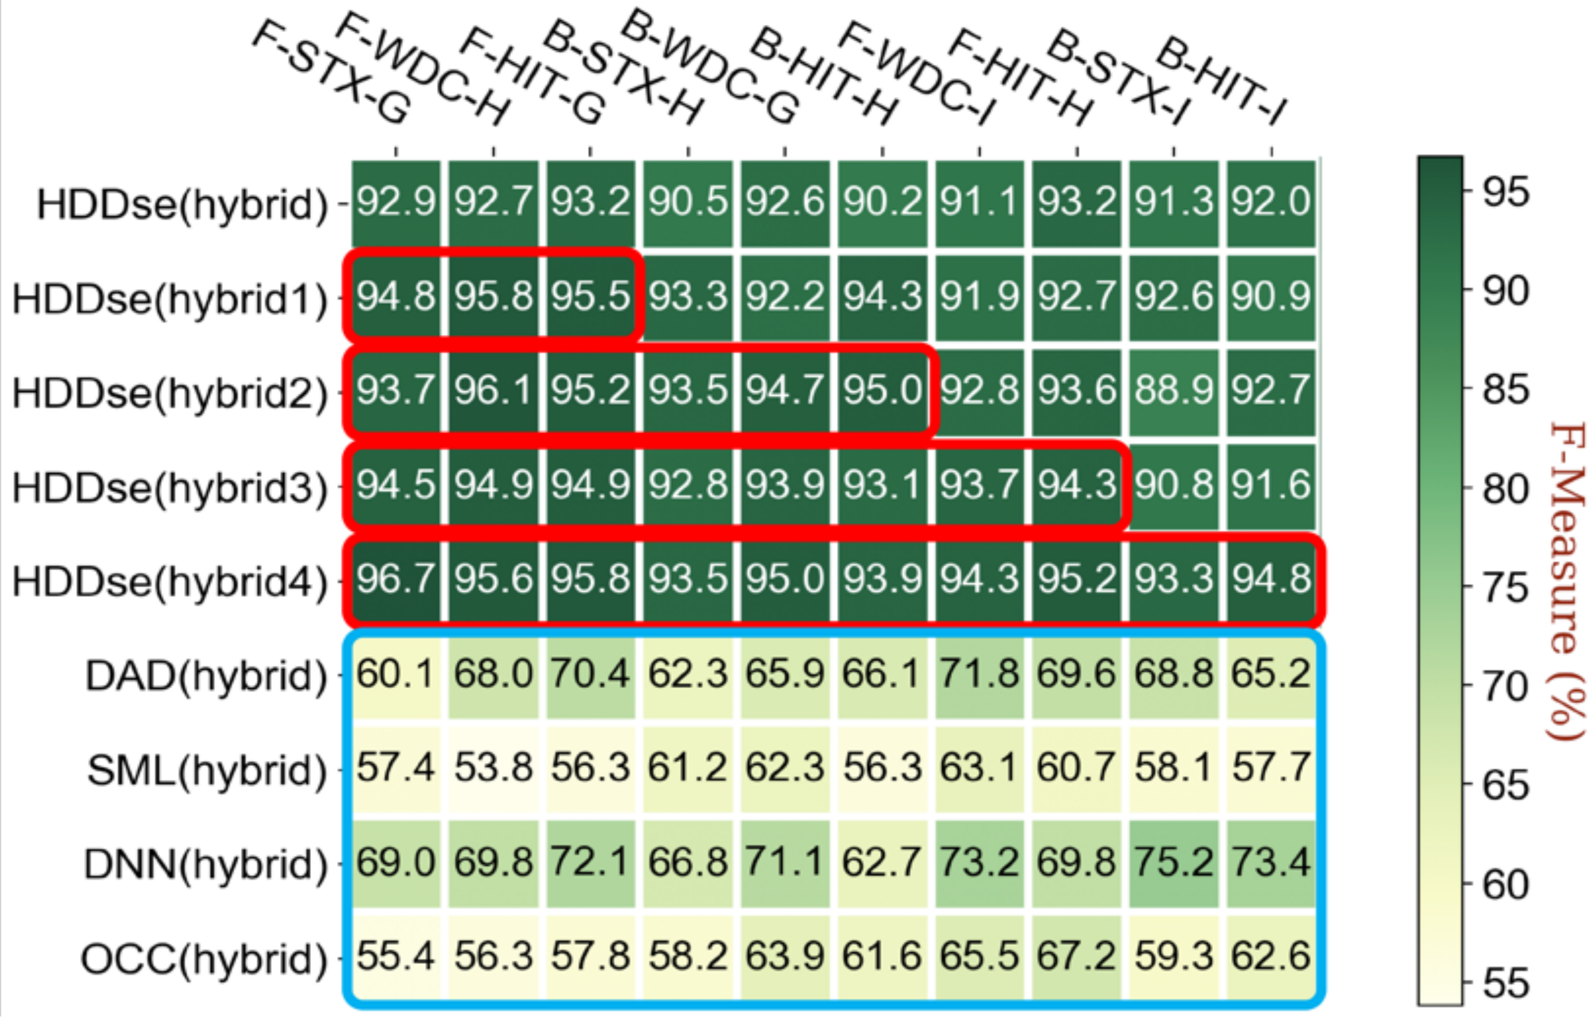

Improvement of Storage System Reliability

Method	$k(TPR)$	**FP**	**FN**	**Cost**	**MTTDL (years)**	**C-MTTDL (hours/dollar)**
OCC	**62.6%**	8212	1062	1,748,600	397.6	1.94
DAD	**45.2%**	3422	1537	838,100	276.7	2.89
SML	**72.6%**	6159	783	1,310,100	504.10	3.37
DNN	**85.3%**	4791	419	1,000,100	814.13	7.13
HDDse	**95.8%**	103	140	34,600	1656.3	419.35

Training and Detecting Time

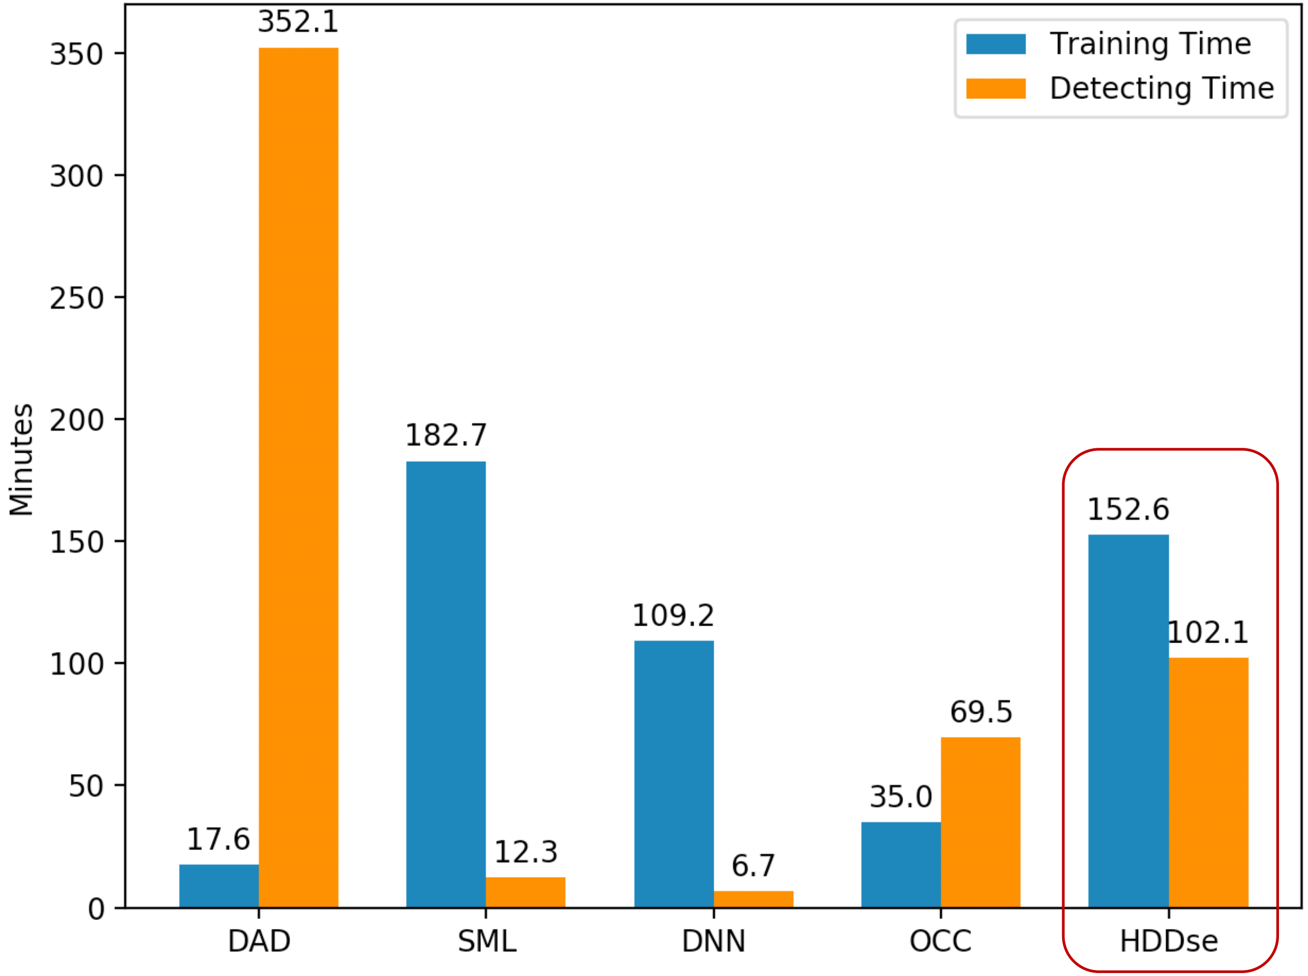

Evaluating Practical Long-Term Availability

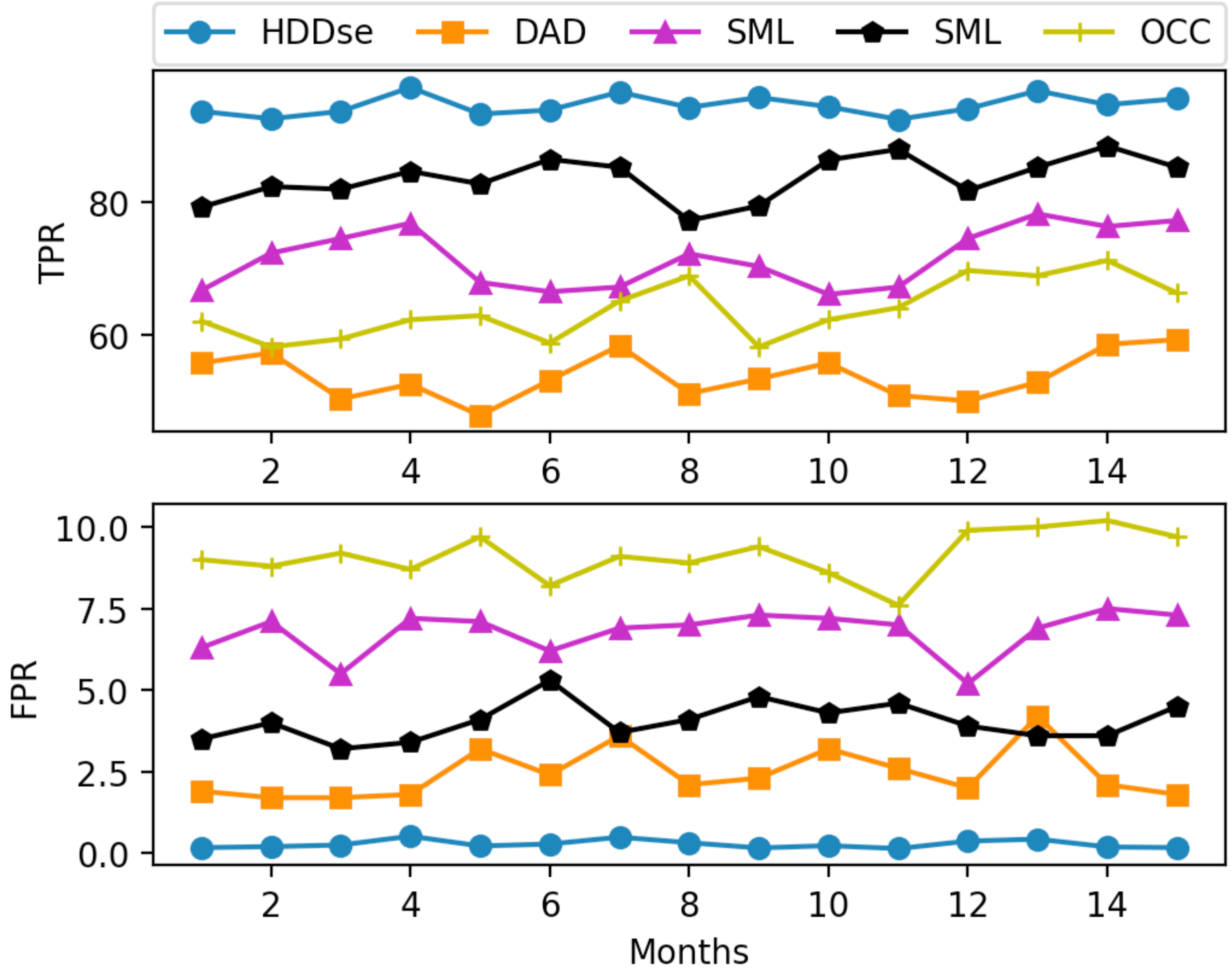

Conclusion

HDDse: an LSTM-based siamese network that can learn the dynamically changed long-term behavior of disk healthy statues and generate a unified and efficient high dimensional disk state embeddings from low dimensional S.M.A.R.T attributes for disk failure detection.

Applicability	Adaptability	Imbalance datasets	Minority Disk	Performance
√	√	√	√	**TPR: 92%-97%** **FPR: 0.2%-0.4%** **F-Measure: 91%-97%**

Thanks