DFPE: Explaining Predictive Models for Disk Failure Prediction

Yanwen Xie, Dan Feng, Fang Wang, Jizhong Han, Xuehai Tang, Xinyan Zhang

ywxie@hust.edu.cn

May 24, 2019

信息存储系统教育部重点实验室
Key Laboratory of Information Storage System
Ministry of Education of China

武汉光电国家研究中心
WUHAN NATIONAL LABORATORY FOR OPTOELECTRONICS

華中科技大學
HUAZHONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Background: Disks

- Disks are widely deployed in datacenters
  - Why?
    - Low cost per bit stored
    - Large storage capacity
    - Mature technology
  - For?
    - Cold data
    - Backup
    - Archiving
    - Long-term
    - et al.

Background: Disk failure prediction

- Disk failures are ordinary events
  - A large number of disks
  - Many disks have been serving for several years
- Advantages of disk failure prediction
  - Keep high reliability
  - Lower the impact of the failure and the overheads of the recover

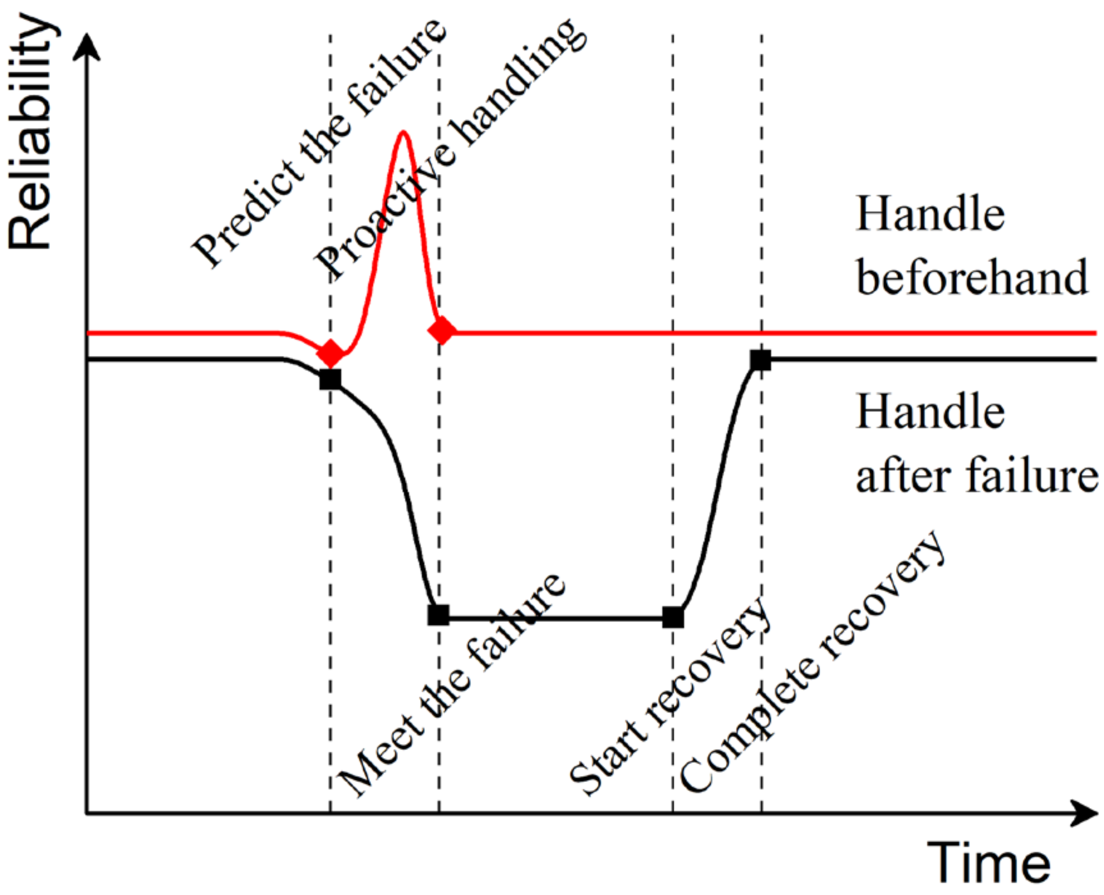

Background: Building a model and deploying

- Collect logs and failures
  - S.M.A.R.T. attributes
- Build machine learning models offline
  - Meet the requirements
    - Fault detection rate (FDR): the higher the better
    - False alarm rate (FAR): the lower the better
- Deploy to online
  - Predict the failure online

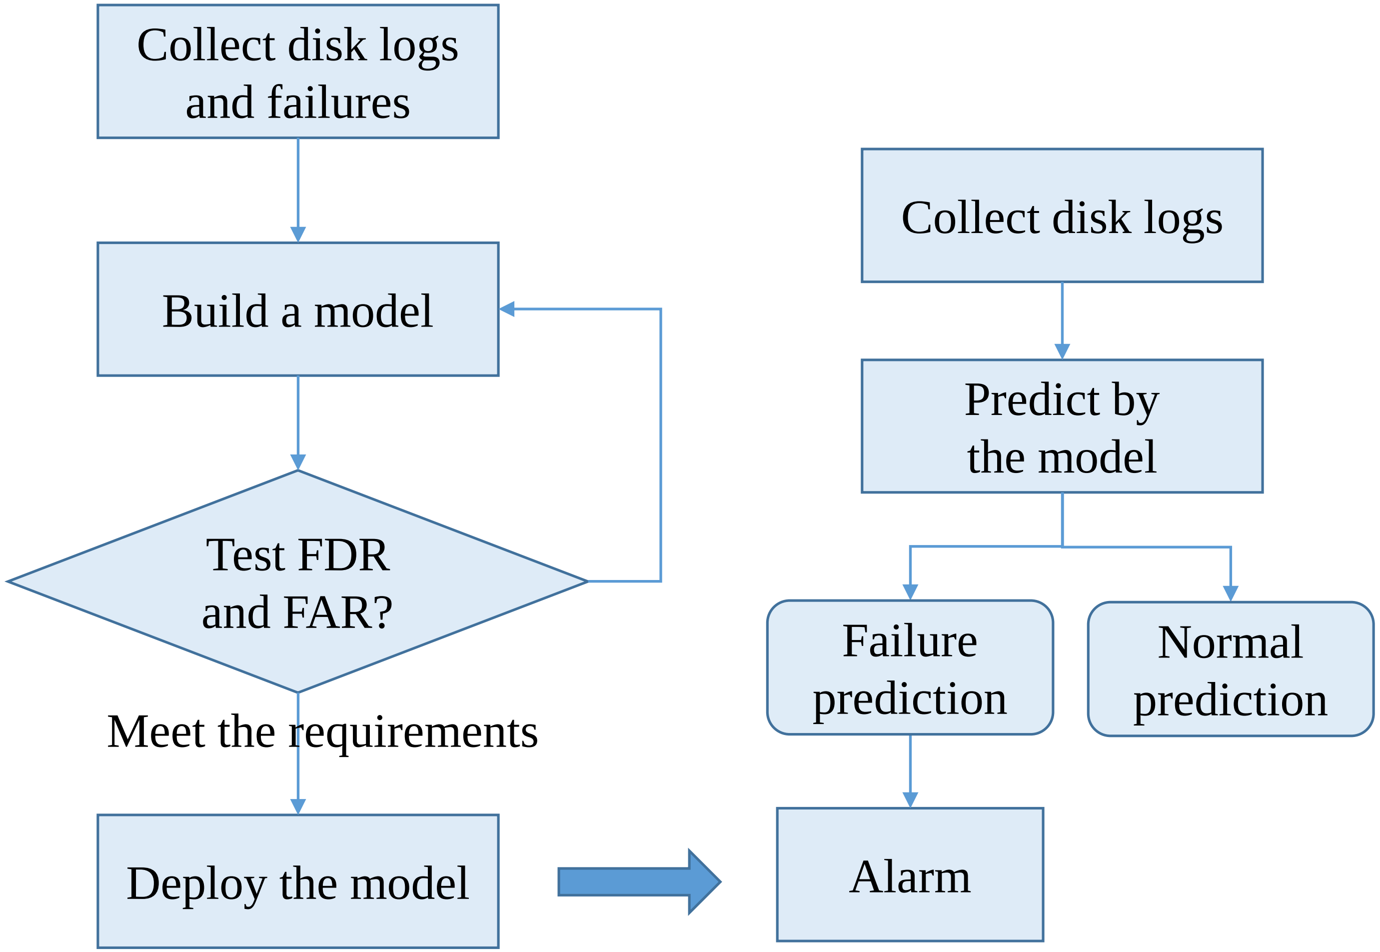

Background: Explainability of a model

- Improve the believability of a model
  - Pass the test: work well for existing cases
  - High explainability: give out reasonable explanations
- Explainability/Interpretability:
  - Not only a result,
  - But also explain how it gets the result from the input
    - Which features are important?
    - How important?
- Advantages of high explainability
  - Expose the bias and over-fitting when unreasonable explanations are presented
  - Improve the believability of a model
  - Output more information related to the result to enable intelligent handling subsequently

Background: Building a model and deploying

- Collect logs and failures
- Build machine learning models offline
  - Meet the requirements
  - Explain the model
- Deploy to online
  - Predict the failure online
  - Explain the failure prediction

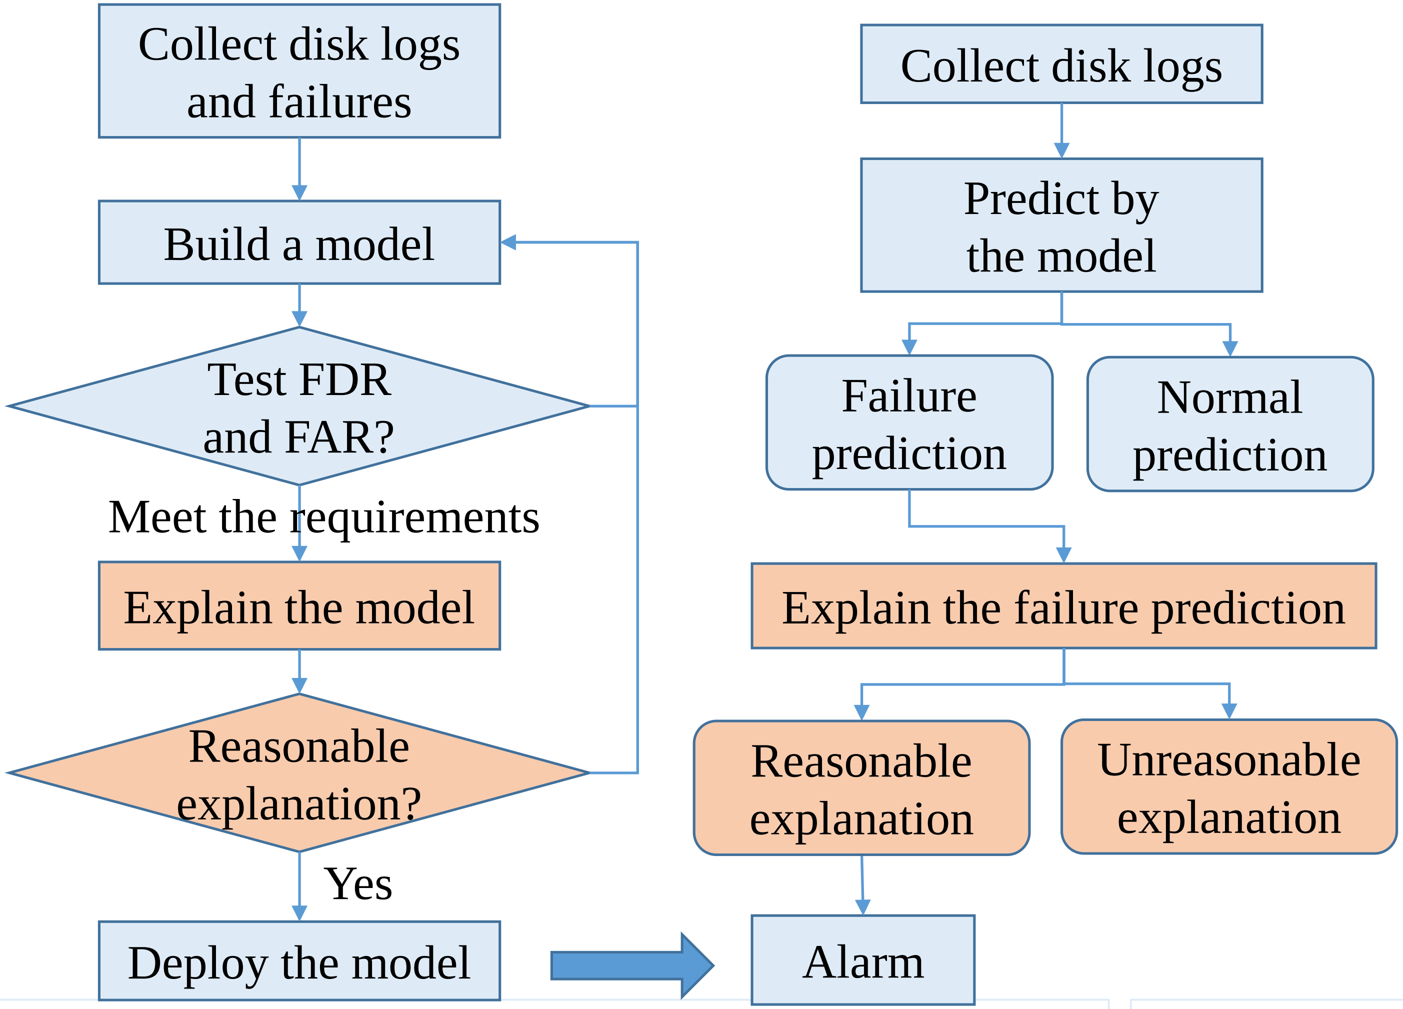

Background: Models for disk failure prediction

- Simple models
  - Decision Trees, Decision Rules, Naïve Bayes, …

Low overheads, Fast, High explainability
Limited learning capacity and accuracy

- Ensemble models
  - Random Forests, GBDT, XGBoost, …

High learning capacity and accuracy
More basic models, lower explainability

- Complex models
  - Neural Networks: MLP, RNN, LSTM, …

High learning capacity and accuracy
More complex, lower explainability

Background: Explanation methods

- Apply explanation methods
  - Keep the high learning capacity and accuracy of the model
  - Improve the explainability of ensemble/complex models
- Global explanation methods
  - Explain the model
  - Like: MDA(Mean Decrease Accuracy), MDI (Mean Decrease Impurity), …
    - Measure the feature importances
- Local explanation methods
  - Explain the output results of the model
  - Like: LIME(KDD’16), …
    - Measure the feature importances

Motivation: The problem

- Complex models are applied to improve the accuracy but with the cost of explainability.
- Current explanation methods can help but the improvement is limited
- Characteristics of disk failure prediction
  - Time series analysis problem
  - Multiple-instance learning problem
    - Unknown failure symptom / time series change point
  - Imbalanced classification problem
    - Failed disks, failure samples and failure predictions are much rarer
    - Only interested in the failure
  - Failure predictions can be caused by multiple causes.

Motivation: The problem

- Global explanation methods (MDA, MDI, …)
  - Not handle the imbalanced
    - The explanation is dominated by the normal disks and normal samples
- Local explanation methods (LIME, …)
  - Not handle the imbalanced and the multiple-instance
    - Extra explanations
- Only the feature importances without considering multiple causes

Motivation: How to solve the problem?

- DFPE: Disk Failure Prediction Explainer
  - Time series analysis problem
    - Support models for time series analysis
  - Multiple-instance learning problem
    - Find the failure symptom / time series change point with the given model
  - Imbalanced classification problem
    - Explain failure predictions ONLY
  - Observe that failure predictions can be caused by multiple causes.
    - Define Minimum Failure Cause Set (MFCS) and find out as many MFCSs as possible

Design: Replacement test

- Only explain failure predictions

$$F_1 \quad F_2 \quad F_3 \quad F_4 \quad F_5 \quad F_6 \quad => \quad Failure$$

- Replacement Test
  - How to omit a feature?
    - Replace the feature values with the mean/median value of the feature of normal disks

$$F_i \quad => \quad \cancel{F_i}$$

  - Failure Cause Set (FCS): omitting the features outside the set does not change the result

$$\left.\begin{array}{l} F_1 \quad F_2 \quad F_3 \quad F_4 \quad F_5 \quad F_6 \quad => \quad Failure \\ \cancel{F_1} \quad \cancel{F_2} \quad F_3 \quad \cancel{F_4} \quad F_5 \quad \cancel{F_6} \quad => \quad Failure \end{array}\right\} \quad FCS = \{F_3, F_5\}$$

Design: Minimum Failure Cause Set (MFCS)

- Minimum Failure Cause Set (MFCS)
  - A MFCS is a FCS
  - No subset of a MFCS is a FCS
- Every feature in a MFCS is essential to support the failure prediction
- A MFCS can be a predictive rule

$$MFCS = \{F_3, F_5\} \Rightarrow$$

Rule: When F_3 and F_5 meet some constraints, the disk would fail in the near future.

Design: Find out MFCSs to explain failure predictions

- Step 1: Test each feature to find out a MFCS
  - Omit each feature and test:
    - Result not changed: Continue
    - Result changed: Add the feature, rollback the feature values, continue

Original :	F_1	F_2	F_3	F_4	F_5	F_6	=>	*Failure*	$MFCS = \{\}$
Omit F_1 :	$\cancel{F_1}$	F_2	F_3	F_4	F_5	F_6	=>	*Failure*	$MFCS = \{\}$
Omit F_2 :	$\cancel{F_1}$	$\cancel{F_2}$	F_3	F_4	F_5	F_6	=>	*Failure*	$MFCS = \{\}$
Omit F_3 :	$\cancel{F_1}$	$\cancel{F_2}$	$\cancel{F_3}$	F_4	F_5	F_6	=>	*Normal*	$MFCS = \{ F_3 \}$
Rollback F_3, Omit F_4 :	$\cancel{F_1}$	$\cancel{F_2}$	F_3	$\cancel{F_4}$	F_5	F_6	=>	*Failure*	$MFCS = \{ F_3 \}$
Omit F_5 :	$\cancel{F_1}$	$\cancel{F_2}$	F_3	$\cancel{F_4}$	$\cancel{F_5}$	F_6	=>	*Normal*	$MFCS = \{ F_3, F_5 \}$
Rollback F_5, Omit F_6 :	$\cancel{F_1}$	$\cancel{F_2}$	F_3	$\cancel{F_4}$	F_5	$\cancel{F_6}$	=>	*Failure*	$MFCS = \{ F_3, F_5 \}$

Design: Find out MFCSs to explain failure predictions

- Continue to find more MFCSs
  - Omit features in found MFCSs and test
    - Case 1: Normal prediction: Done

$$F_1 \quad F_2 \quad \cancel{F_3} \quad F_4 \quad \cancel{F_5} \quad F_6 \quad => \quad Normal$$

    - Case 2: Failure prediction: Go on a new round of tests

$$F_1 \quad F_2 \quad \cancel{F_3} \quad F_4 \quad \cancel{F_5} \quad F_6 \quad => \quad Failure$$

$$\cancel{F_1} \quad \cancel{F_2} \quad \cancel{F_3} \quad F_4 \quad \cancel{F_5} \quad \cancel{F_6} \quad => \quad Failure \qquad MFCS_2 = \{ F_4 \}$$

Design: Find out MFCSs to explain failure predictions

- Can only find out MFCSs without common features
  - e.g. $\{ F_3 , F_5 \}$ and $\{ F_4 \}$
    - First round

$$\left.\begin{array}{lllllll} F_1 & F_2 & F_3 & F_4 & F_5 & F_6 & \Rightarrow \text{Failure} \\ \cancel{F_1} & \cancel{F_2} & F_3 & \cancel{F_4} & F_5 & \cancel{F_6} & \Rightarrow \text{Failure} \end{array}\right\} \quad MFCS = \{ F_3, F_5 \}$$

    - Second round

$$\left.\begin{array}{lllllll} F_1 & F_2 & \cancel{F_3} & F_4 & \cancel{F_5} & F_6 & \Rightarrow \text{Failure} \\ \cancel{F_1} & \cancel{F_2} & \cancel{F_3} & F_4 & \cancel{F_5} & \cancel{F_6} & \Rightarrow \text{Failure} \end{array}\right\} \quad MFCS_2 = \{ F_4 \}$$

  - What if the second MFCS is $\{ F_2 , F_3 \}$?
    - F_3 has been omitted after is $\{ F_3, F_5 \}$ found

Design: Find out MFCSs to explain failure predictions

- Step 2: Validate known MFCSs (Optional)
  - Known MFCSs: the MFCSs found from the explanations for existing failure predictions.
    - $\{km_1, km_2, km_3, km_4, \ldots\}$
  - Found MFCSs from the previous step
    - $\{fm_1, fm_2, fm_3, fm_4, \ldots\}$
  - Filters to reduce the validation times
    - Not subset

      $\exists fm_j: km_i \subset fm_j \quad \Rightarrow \quad km_i$ is not a FCS $\quad \Rightarrow \quad km_i$ is not a MFCS
    - Not superset

      $\exists fm_j: fm_j \subset km_i \quad \Rightarrow \quad km_i$ is not the minimal $\quad \Rightarrow \quad km_i$ is not a MFCS
    - Should have common features

      $\exists fm_j: fm_j \cap km_i \neq \emptyset$

Design: Measure feature importances on MFCSs

- Calculate feature importances for each MFCS individually
  - Measure how much a feature has to be changed to change the prediction.
  - e.g. $MFCS = \{ F_3, F_5 \}$, to measure the importance of F_3

$$\cancel{F_1} \quad \cancel{F_2} \quad F_3 \quad \cancel{F_4} \quad F_5 \quad \cancel{F_6} \quad => \quad Failure$$

$$\cancel{F_1} \quad \cancel{F_2} \quad \cancel{F_3} \quad \cancel{F_4} \quad F_5 \quad \cancel{F_6} \quad => \quad Normal$$

  - Find out the change point $\widehat{F_3}$ between F_3 and $\cancel{F_3}$ with binary search
  - Measure the feature importance

$$IMP_{F_3, MFCS} = \frac{|F_3 - \widehat{F_3}|}{|F_3 - \cancel{F_3}|}$$

Design: Gather explanations to explain models

- Calculate metrics for each known MFCS
  - TP_{MFCS}: the number of failed disks predicted successfully with the MFCS
  - FP_{MFCS}: the number of normal disks predicted to fail with the MFCS
  - Detection Rate: the importance/popularity of the predictive rule

$$FDR_{MFCS} = \frac{TP_{MFCS}}{\text{the } number\ of\ failed\ disks}$$

  - False Alarm Rate: the believability of the predictive rule

$$FAR_{MFCS} = \frac{FP_{MFCS}}{\text{the } number\ of\ normal\ disks}$$

Design: Gather explanations to explain models

- Measure feature importances on a model
  - TP_{F_i}: the number of failed disks predicted successfully with any MFCS including F_i
  - Feature Importance:

$$IMP_{F_\mathrm{i}} = \frac{TP_{F_i}}{the\ number\ of\ failed\ disks}$$

Evaluation: Datasets

- Datasets

Label	Disk Series	Collected From	Download	Normal disks	Failed disks	Sampling Interval	Total time
D0	Seagate ST31000524NS	Baidu Company	[6]	22962	433	1 hour	1 week or 20 days[1]
D1	Seagate ST4000DM000	Backblaze Company	[26]	34295	2502	1 day	Feb 2014 $\sim$ Sep 2017
D2	Seagate ST3000DM001			2898	1006		Feb 2014 $\sim$ Nov 2015
D3	Seagate ST31500541AS			1679	238		Feb 2014 $\sim$ Sep 2017
D4	Hitachi HDS722020ALA330			4535	193		Feb 2014 $\sim$ Sep 2017
D5	WDC WD30EFRX			1161	152		Feb 2014 $\sim$ Sep 2017
D6	HGST HMS5C4040ALE640			8569	126		Feb 2014 $\sim$ Sep 2017
D7	HGST HMS5C4040BLE640			16181	120		Mar 2014 $\sim$ Sep 2017
D8	Hitachi HDS5C3030ALA630			4512	116		Feb 2014 $\sim$ Sep 2017
D9	Seagate ST8000DM002			9882	110		May 2016 $\sim$ Sep 2017

[1] For D0, 1 week samples are collected for normal disks and 20 days before the failure for failed disks

- Visualized explanations on a Random Forests model on D0
- Overheads on the 10 datasets.

Evaluation: Procedure

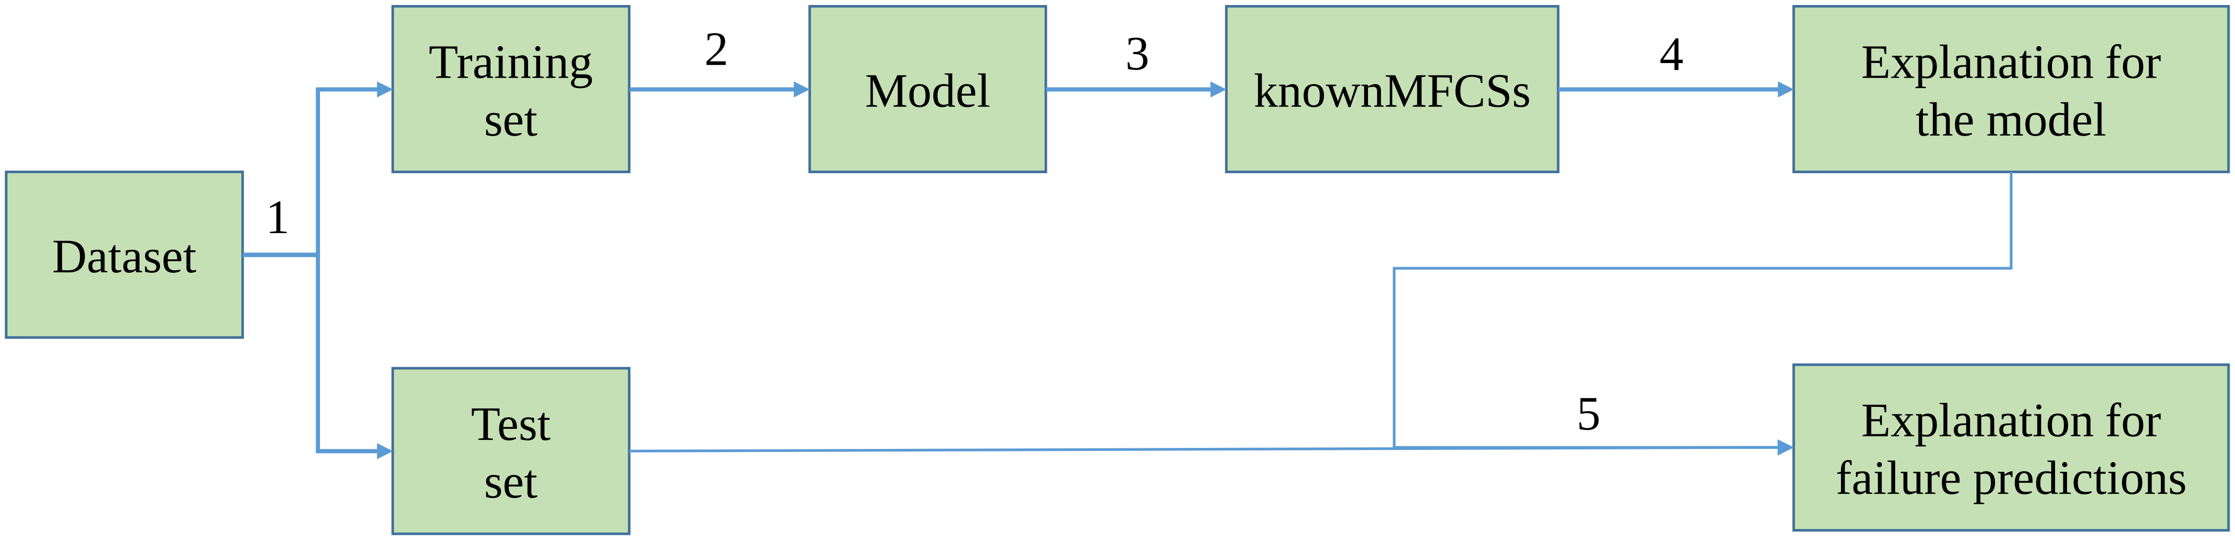

1: Split the dataset

2: Train a predictive model

3: Build knownMFCSs with Step 1

4: Perform Step 2 on training data and explain the model

5. Explain the failure predictions in test set

Evaluation: Visual explanation for models

Inferred Rule ($MFCS$)	$FDR_{MFCS} \downarrow$ [1]	FAR_{MFCS}
$\{5\}$	0.7822	0.00019
$\{8\}$	0.3201	0.00075
$\{3, 9, 11\}$	0.3168	0.00137
$\{6\}$	0.1518	0.00000
$\{8, 9, 11\}$	0.0495	0.00006
$\{1, 3, 6\}$	0.0462	0.00000
$\{6, 11\}$	0.0396	0.00000
$\{3, 5\}$	0.0330	0.00093
$\{4, 8, 12\}$	0.0330	0.00006
$\{2, 7, 8, 11\}$	0.0330	0.00006
...	...	...

Inferred Rule ($MFCS$)	FDR_{MFCS}	$FAR_{MFCS} \downarrow$ [1]
$\{3, 9, 11\}$	0.3168	0.00137
$\{3, 5\}$	0.0330	0.00093
$\{8\}$	0.3201	0.00075
$\{5\}$	0.7822	0.00019
$\{3, 4, 8, 9\}$	0.0198	0.00012
$\{1, 8\}$	0.0033	0.00006
$\{4, 8, 12\}$	0.0330	0.00006
$\{6, 8, 11\}$	0.0231	0.00006
$\{8, 9, 11\}$	0.0495	0.00006
$\{4, 8, 11\}$	0.0099	0.00006
...	...	...

- MFCSs sorted by detection rates
- MFCSs sorted by false alarm rates
- Feature importances

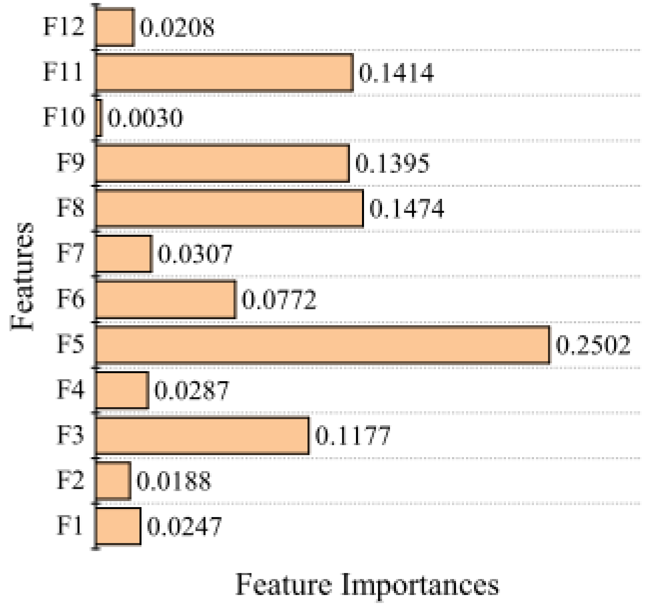

Evaluation: Explanations from MDA and MDI

- Feature importances

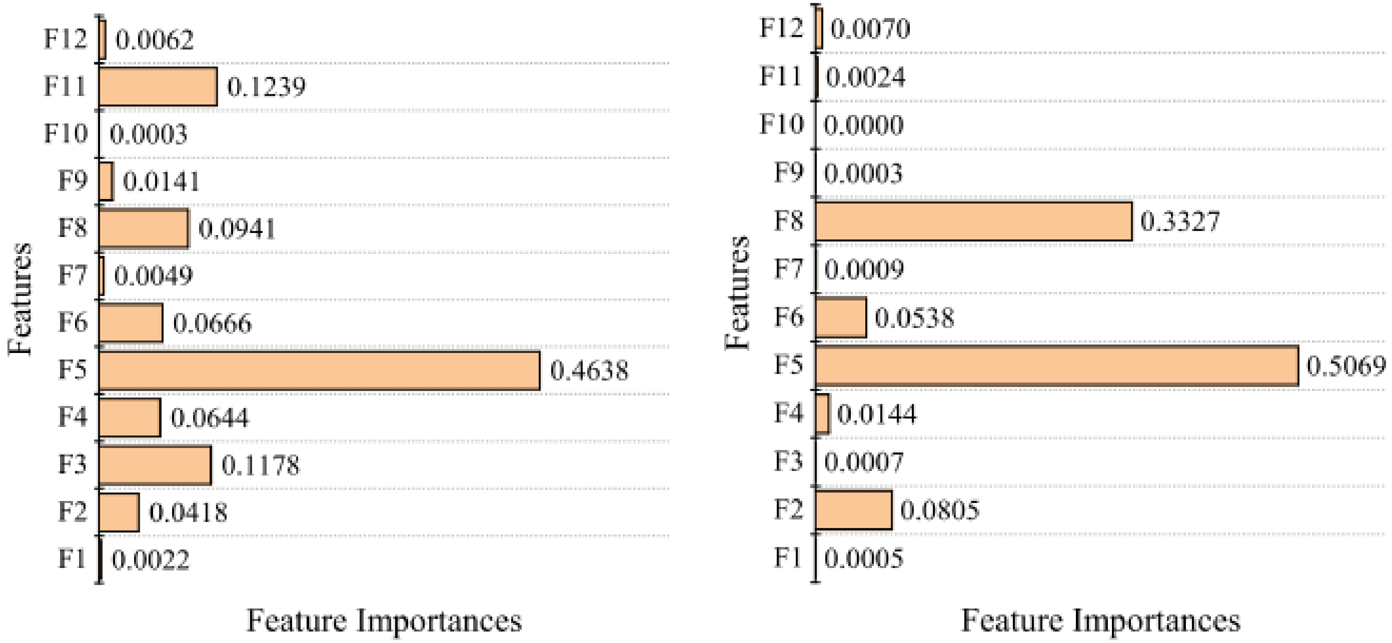

(a) MDI

(b) MDA

Evaluation: Comparison

- DFPE Explains more
  - Predictive rules with their detection rates and false alarm rates
  - Can be applied in post-process
    - Remove rules with a low detection rate or a high false alarm rate to improve a model

Inferred Rule ($MFCS$)	$FDR_{MFCS} \downarrow$ [1]	FAR_{MFCS}
$\{5\}$	0.7822	0.00019
$\{8\}$	0.3201	0.00075
$\{3, 9, 11\}$	0.3168	0.00137
$\{6\}$	0.1518	0.00000
$\{8, 9, 11\}$	0.0495	0.00006
$\{1, 3, 6\}$	0.0462	0.00000
$\{6, 11\}$	0.0396	0.00000
$\{3, 5\}$	0.0330	0.00093
$\{4, 8, 12\}$	0.0330	0.00006
$\{2, 7, 8, 11\}$	0.0330	0.00006
...	...	...

Inferred Rule ($MFCS$)	FDR_{MFCS}	$FAR_{MFCS} \downarrow$ [1]
$\{3, 9, 11\}$	0.3168	0.00137
$\{3, 5\}$	0.0330	0.00093
$\{8\}$	0.3201	0.00075
$\{5\}$	0.7822	0.00019
$\{3, 4, 8, 9\}$	0.0198	0.00012
$\{1, 8\}$	0.0033	0.00006
$\{4, 8, 12\}$	0.0330	0.00006
$\{6, 8, 11\}$	0.0231	0.00006
$\{8, 9, 11\}$	0.0495	0.00006
$\{4, 8, 11\}$	0.0099	0.00006
...	...	...

Evaluation: Comparison

- Feature importances
  - Distribute more evenly: easier for feature comparison
  - Better for the imbalanced learning problem

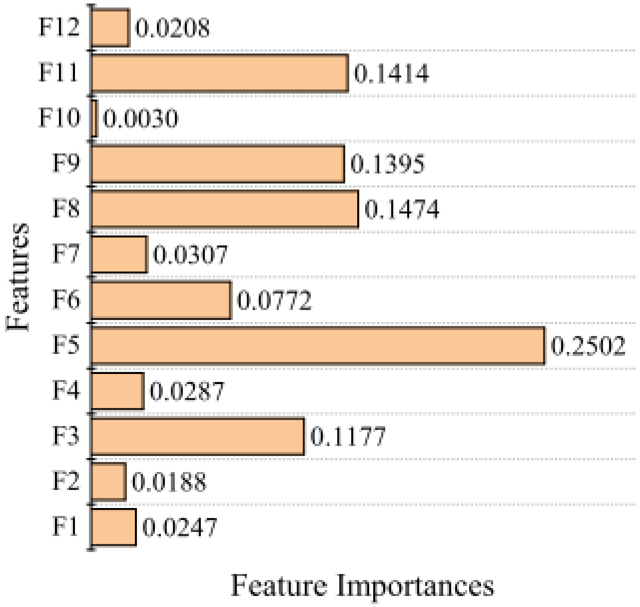

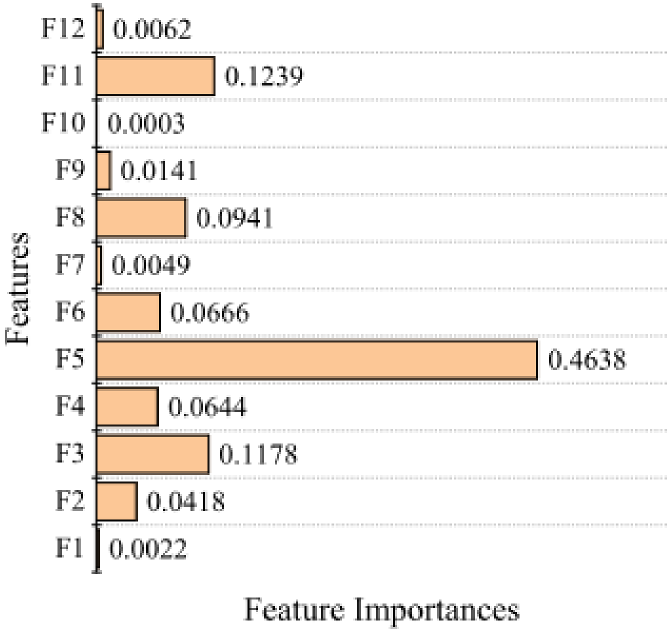

(a) MDI

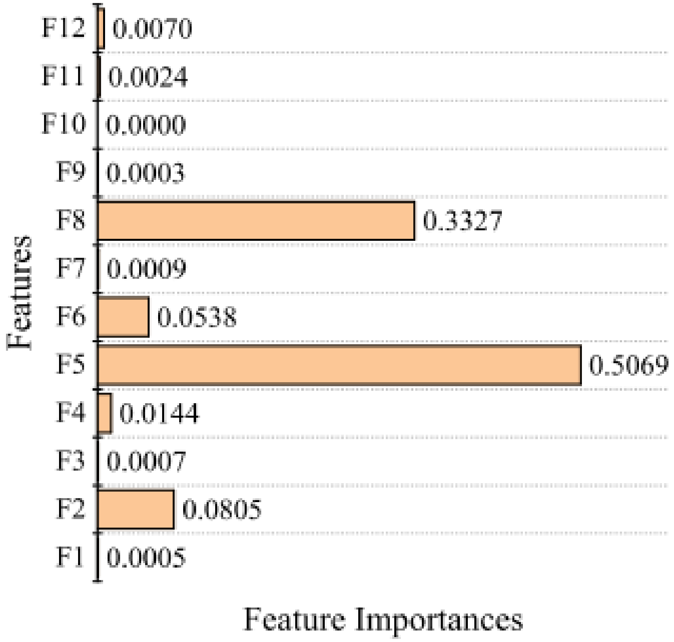

(b) MDA

Evaluation: Visual explanation for failure predictions

- Found MFCSs
- Detection rates and false alarm rates for the MFCSs
- Feature importances on the MFCSs

$MFCS$	FDR_{MFCS}	FAR_{MFCS}
$\{6\}$	0.1518	0.00000
$\{5\}$	0.7822	0.00019
$\{3, 9, 11\}$	0.3168	0.00137

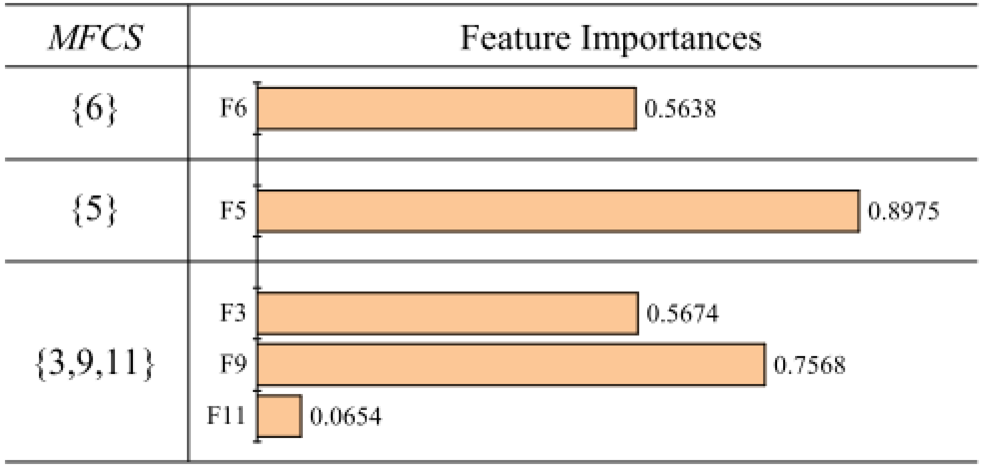

Evaluation: Explanation from LIME

- Feature importances

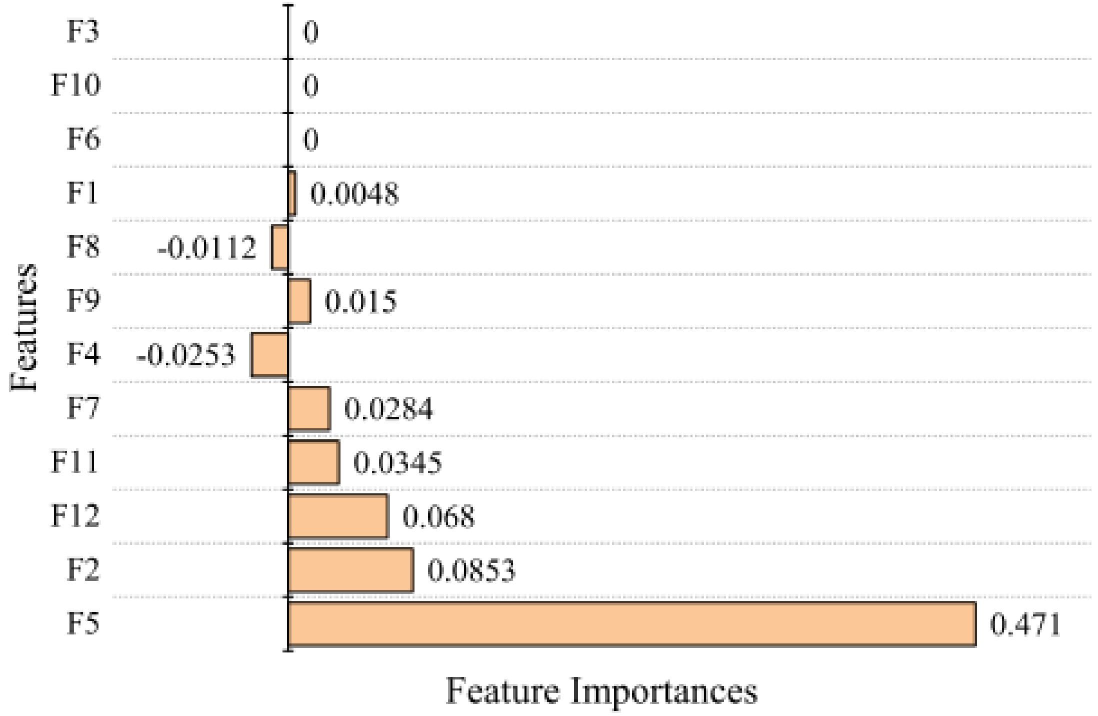

Evaluation: Comparison

- DFPE can
  - Find out multiple causes
  - Measure the feature importances individually
  - Provides detection rates and false alarm rates

$MFCS$	FDR_{MFCS}	FAR_{MFCS}
$\{6\}$	0.1518	0.00000
$\{5\}$	0.7822	0.00019
$\{3, 9, 11\}$	0.3168	0.00137

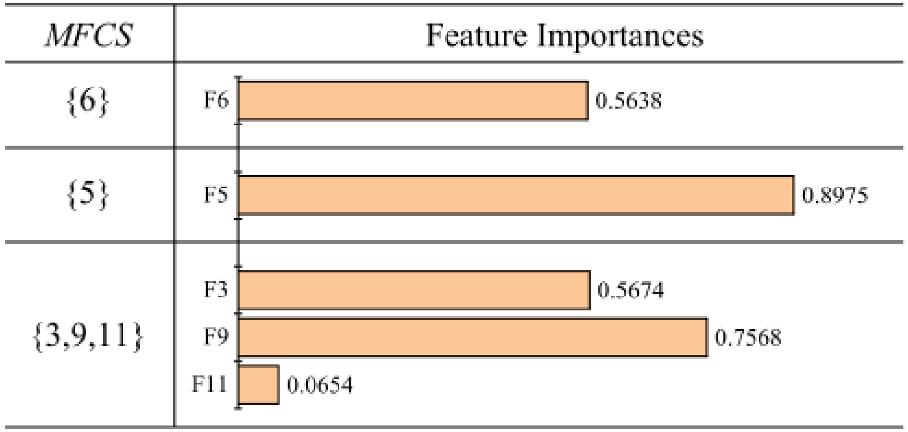

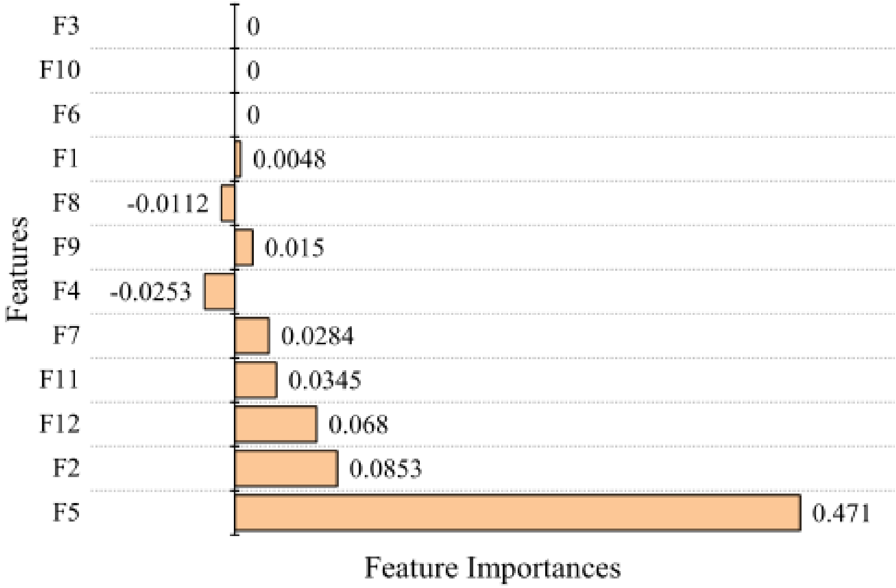

Evaluation: Detect the hidden bias

- F_5: the power-on time
- *F_5 is important* observed from the explanations of MDA, MDI and LIME

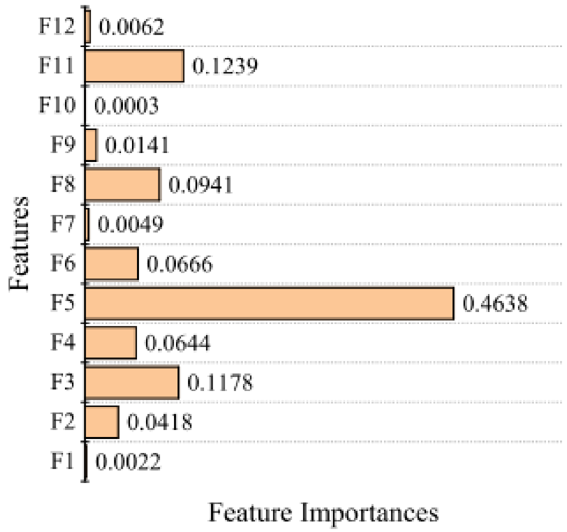

(a) MDI

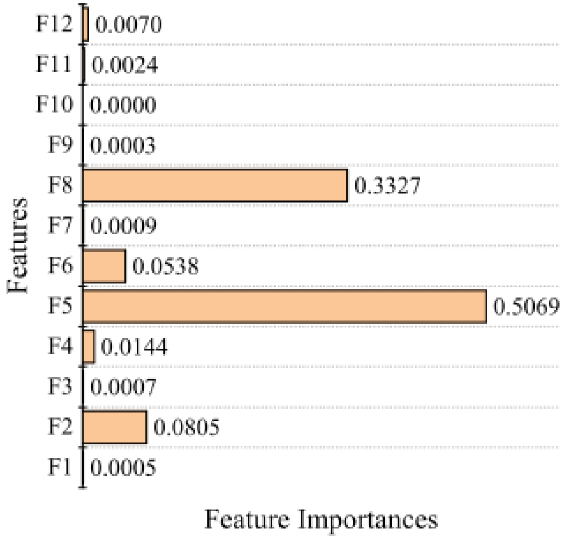

(b) MDA

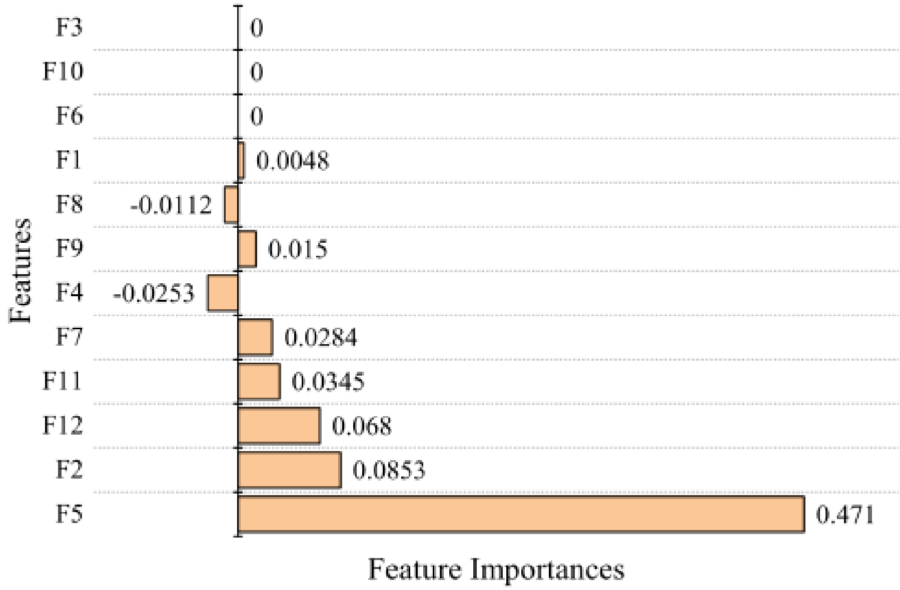

Evaluation: Detect the hidden bias

- *F_5 is* a determining factor observed from the explanation of DFPE

Inferred Rule ($MFCS$)	$FDR_{MFCS} \downarrow$ [1]	FAR_{MFCS}
$\{5\}$	0.7822	0.00019
$\{8\}$	0.3201	0.00075
$\{3, 9, 11\}$	0.3168	0.00137
$\{6\}$	0.1518	0.00000

$MFCS$	FDR_{MFCS}	FAR_{MFCS}
$\{6\}$	0.1518	0.00000
$\{5\}$	0.7822	0.00019
$\{3, 9, 11\}$	0.3168	0.00137

- The predictive rule: When the power-on time of a disk exceeds a threshold, the disk will fail.
  - Age bias?

Evaluation: Detect the hidden bias

- The values are normalized to [-1, 1]
  - The smaller, the more power-on time
- The bias is caused by the data bias in the dataset.
  - Normal disks: [-0.08, 1]
  - Failed disks: [-1, 1]
  - No normal samples with large power-on time values
- Maybe the dataset was collected by:
  - Export records of all disks at a time
  - Update records of failed disks afterwards

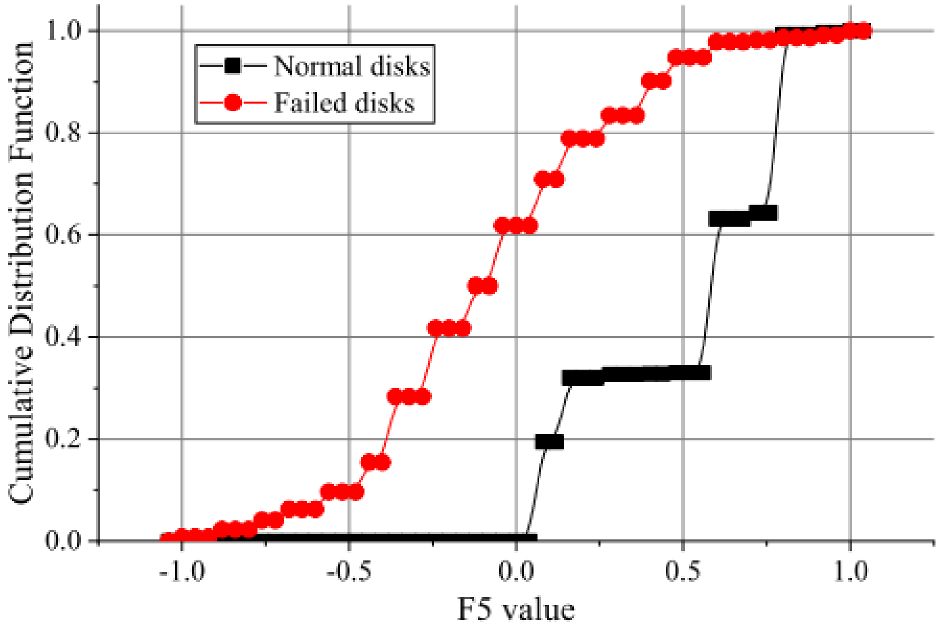

Evaluation: Handle the hidden bias

- Two methods:
  - Predict a disk to fail only when there are other MFCSs except $\{F_5\}$ in its explanation.
  - Remove F_5 and rebuild the model
- After handling the hidden bias
  - Prediction accuracy on the test set decreases
  - More applicable because it does not have the unsound rule.

Method	FDR	FAR
Original	0.8769	0.0033
1	0.6385	0.0029
2	0.6769	0.0109

Evaluation: The overheads

- DFPE needs more time than MDA and LIME, but
  - Compared to the overheads of failure handling
  - Given the advantages of high explainability
- The extra overheads are acceptable

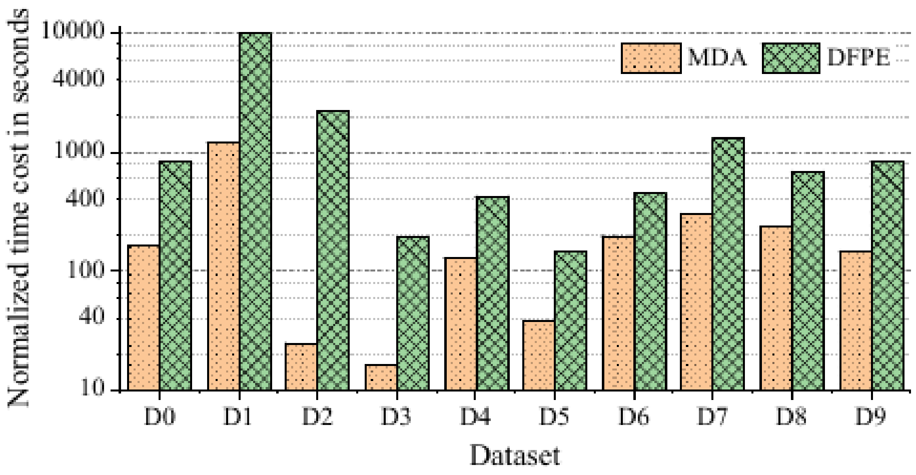

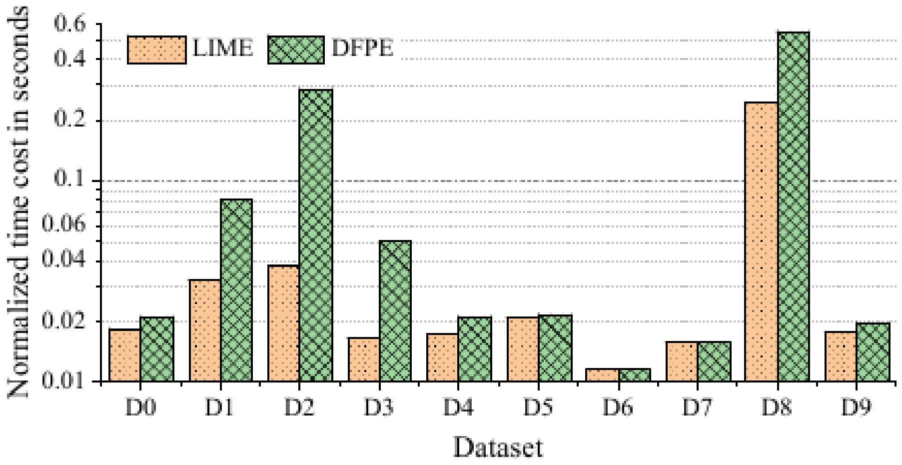

Summary

- Emphasize the importance of explainability
- Point out the data bias in a popular dataset
- Propose an explanation method for complex models in disk failure prediction
- Present a case on how the new method helps to detect and handle bias
- Provide a new perspective of measuring feature importances
- Enable intelligent failure handling by providing the failure causes.

- Future work
  - Seek more applications for the new explanation method
  - Lower the overheads

Q&A

- Thank you very much!