



Proactive Failure Detection of Hard Disk Drives

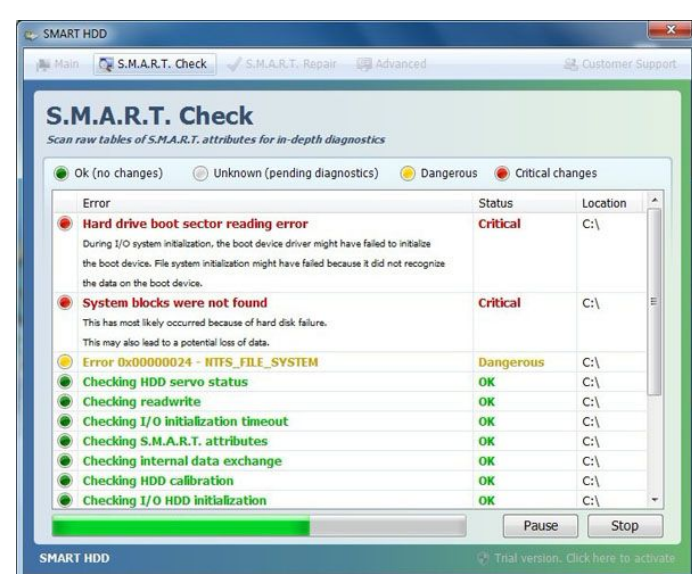
Wendy Li (liwendy) (CS221), Ivan Suarez (isuarezr) (CS221), Juan Camacho (jcamach2) (CS229)



Abstract

Datacenter downtime costs are increasing significantly in the past years from \$5,600/minute in 2010 to \$8,851/minute in 2016. Hard disk drives (HDD) are among the most common failing components in data centers. Today, HDD failure prevention uses S.M.A.R.T. attributes to predict disk failure through simple value thresholding. This results in high false positive rates which makes replacement of healthy disks very expensive. **The goal of this project is to apply machine learning techniques to accurately and proactively predict hard disk drive failures.** In evaluating what method produces the highest accuracy model, we implemented Logistic Regression, Random Forest and Naïve Bayes algorithms.

S.M.A.R.T. Features



Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) is a monitoring system included in all computer hard disk drives and used to detect and report on various indicators of drive reliability.

Figure 1. Typical SMART Check program

Drive Status	SMART 5	SMART 187	SMART 188	SMART 197	SMART 198
Operational	1.1%	0.5%	4.8%	0.7%	0.3%
Failed	42.2%	43.5%	44.8%	43.1%	33.0%

Figure 2. BackBlaze correlation percentages between SMART features and hard disk drive health

S.M.A.R.T. 5 - Reallocated Sectors Count

When the drive's logic believe that a sector is damaged, it can remap the faulty sector number to a new physical sector drawn from a pool of spares.

S.M.A.R.T. 187- Reported Uncorrectable Errors

The count of errors that could not be recovered using hardware ECC. Large scan error counts can be indicative of surface defects and therefore are believed to be indicative of lower reliability.

S.M.A.R.T. 188- Command Timeout

The count of aborted operations due to HDD timeout.

S.M.A.R.T. 197- Probational Count

Disk drives put suspect bad sectors on probation until they either fail permanently and are reallocated or continue to work without problems.


S.M.A.R.T. 198 - Offline Uncorrectable Sector Count

The total count of uncorrectable errors when reading/writing to a sector. A rise in the value of this attribute indicated defects of the disk surface and/or problems in the mechanical subsystem.

Model Selection

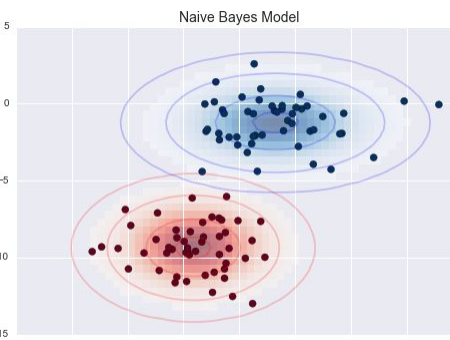
Logistic Regression

Logistic regression is a linear regression technique that can be used to predict binary-class instances.



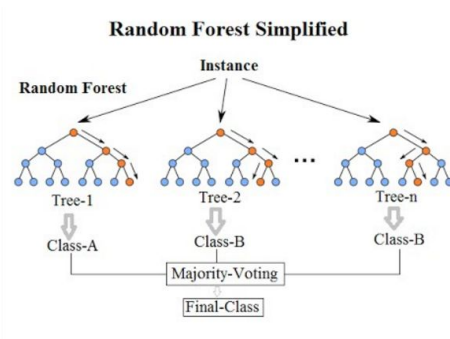
Naïve Bayes Classifier

Naïve Bayes classifier is a simple learning algorithm based on Baye's theorem. The learning phase analyzes the dataset and builds probability distribution models of the attributes. When the models are obtained, the prediction is carried out by calculating the probability of all attributes under the assumption that all attributes are independent and identically distributed.



Random Forest

Random forest is a collection of decision trees. This technique can be viewed as meta-learning, which improves the prediction quality by casting votes among the trees and assigning the most voted class to the predicted instance.



Data

All data is collected from the S.M.A.R.T. readings of BackBlaze's storage hard disks.

Training Data Set:

HDD data from Q1 2017 and Q2 2017
• Q1: 85,301 disks Q2; 89838 disks
• Total failed drives: 778

Test Data Set:

HDD data from Q1 2016
• Total disks: 64,074
• Total failed disks: 337

Features used: raw and normalized data from SMART 5, 187, 188, 197, 198 . Failure/ operational status, data, serial number

Failure definition: A drive is considered to have failed if it was replaced as part of a repairs procedure.

Preprocessing

- Convert S.M.A.R.T. features into:
 - normalized
 - raw
 - binary (equal to 0, greater than 0) data points
- If a failure occurs → mark last 60 days of the hard disk drive's failure/operational status as failed
- Option 1: Balance the data set → # of failed data points = # of operational data points
- Option 2: Balance the data set → Weight on of failed data point = # of operational disks/ # of failed disks , Weight on operation disk data points = 1

Approach

- Preprocess data and propagate failures back n days
- Split data into 66% train, 33% test
- Based on highest F-measure choose between a balanced 50/50 or untouched data set
- Compare F-measure of results with data that is raw, normalized and binary
- Perform error analysis on model with highest F-measure and determine other methods to tune performance
- Evaluate data based on complete unbalanced data set

Optimal Data Set Types for each Algorithm

Algorithm	Data preprocessing
Logistic Regression	Unbalanced data set + binary SMART features
Naïve Bayes	Unbalanced data set + raw SMART features
Random Forest	Unbalanced data set + normalized SMART features + 4 depth

Error Analysis

Analyzing False Positives

Observations:

For normalized values, the closer to 255, the healthier the disk. Logistic regression tends to classify disks with low SMART 187, 188 as about to fail when in fact it doesn't.

Suggestion: decrease the weight of 187, 188.

SMART 5,187, 188, 197, 198
100, 100, 100, 100, 100
100, 61, 100, 100, 100
78, 0, 0, 100, 100
100, 64, 100, 100, 100
31, 0, 0, 100, 100

Analyzing False Negatives

Observations:

Logistic Regression does a fairly good job at predicting hard disks that seem healthy as continuing to be operational.

Unfortunately, not all hard disks that fail show signs of failure. For hard disks that look healthy and show no signs of errors, it is nearly impossible to predict failure.
Suggestion: incorporate more features that have weaker correlations to failure than the current 5 features.

SMART 5,187, 188, 197, 198
100, 96, 100, 100, 100
100, 100, 100, 100, 100
100, 100, 100, 100, 100
100, 99, 100, 100, 100
100, 96, 100, 100, 100

Results

Performance Metrics

$$\text{Sensitivity} = \frac{TP}{TP + FN}$$

$$\text{False Positive Rate} = \frac{FP}{TN + FP}$$

$$\text{Specificity} = \frac{TN}{TN + FP}$$

$$\text{False Discovery Rate} = \frac{FP}{TP + FP}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{F1 Score} = \frac{2TP}{2TP + FP + FN}$$

	Good drive	Failing drive
Predicted as good	True negative (TN)	False negative (FN)
Predicted as failing	False positive (FP)	True positive (TP)

Final Results

	Random Forest	Baseline	Logistic Regression	Naive Bayes
Precision = TP/(TP +FP)	0.714	0.0388	0.561	0.167
Recall = TP/(TP+FN)	0.0297	0.722	0.0682	0.513
Accuracy = TP+TN/(TP+TN+FP+FN)	0.995	0.943	0.995	0.984
F-measure = 2*TP/(2TP + FP+FN)	0.057	0.0737	0.122	0.252
TP rate = TP/(TP + FN)	29.7	72.2	6.82	51.3
FP rate = FP/(FP+TN)	0.63	5.64	0.028	1.4
FN rate = FN/(FN+TP)	97	27.8	93.2	48.7
TN rate = TN/(TN+FP)	99.9	94.4	99.9	98.6

Analysis

Logistic Regression – LR performs poorly due to large amounts of “noise” in the data set. Disks with the same feature values can be failing and operational at the same time. LR is unable to easily distinguish between the two.

Naïve Bayes – NB is able to combine the contributions of multiple predictors that in isolation have low predictive power. Naive Bayes classifiers are also robust in the face of irrelevant attributes. Of our current models, a Naïve Bayes Classifier is best suited to our data set type.

Random Forest – Single decision trees are difficult to use, as each node in the tree makes a decision based on a single feature, and furthermore does so only in a binary fashion (i.e., is feature x > 50?). Being not so good predictors, combining the results of multiple trees may not necessarily yield a better predictor.

References

Goldszmidt, Moises. "Finding Soon-to-Fail Disks in a Haystack." *HotStorage*. 2012.

Hamerly, Greg, and Charles Elkan. "Bayesian approaches to failure prediction for disk drives." *ICML*. Vol. 1. 2001.

Murray, Joseph F., Gordon F. Hughes, and Kenneth Kreutz-Delgado. "Machine learning methods for predicting failures in hard drives: A multiple-instance application." *Journal of Machine Learning Research* 6.May (2005): 783-816.

Pinheiro, Eduardo, Wolf-Dietrich Weber, and Luiz André Barroso. "Failure Trends in a Large Disk Drive Population." *FAST*. Vol. 7. No. 1. 2007.

Pitakrat, Teerat, André van Hoorn, and Lars Grunke. "A comparison of machine learning algorithms for proactive hard disk drive failure detection." *Proceedings of the 4th international ACM Sigsoft symposium on Architecting critical systems*. ACM, 2013.