Statistical Analysis of Hard Disk Drive Failure

Executive Summary

Philip Matthew Unrue

Western Governors University

The Problem and Hypothesis

The failure of hard disk drives is inevitable, causing one of the most impactful problems that businesses can experience today as simple drive recovery can cost up to $7,500 per drive (Painchaud, 2018), assuming the data can be recovered. For data centers, keeping multitudes of businesses' data intact for their own operations is crucial. Being able to predict which hard drives are at the highest risk of failure based on understanding of combinations of routine diagnostics test results is an ideal solution to support the backup and replacement of failing drives before the data is lost. Hypothesizing that the study factors, SMART attribute values obtained from daily diagnostics recordings, are not statistically significant in predicting failure, this project shows that evidence exists to reject the null hypothesis.

The Data Analysis Process

The 4th quarter daily data (Backblaze, 2020) of hard drive attributes was combined into a single dataset and then loaded into pandas for tidying and analysis, resulting in a dataframe of nearly 11 million rows representing hard drive days by 131 columns. 678 of these rows represented HDD failures, resulting in an imbalance ratio (Weiss, 2013) of approximately 16,210:1. The rarity of the positive failure cases was the reason that the entire 4th quarter dataset was required.

63 redundant columns, representations of the normalized SMART attribute values were removed to free memory and space for custom calculations. The 63 SMART attribute raw data columns all had NaN values. 10 of the columns had NaN row values exclusively and were dropped. The model column was used to determine specific model and the manufacturer of each drive. A few thousand rows were removed as a result of error values in important grouping

columns like capacity, and a couple hundred rows removed due to representation non-standard hard disk drives like RAID array controllers and generic error values for their model. Two SMART attribute columns were combined as they represented the same information.

The SMART attribute columns with at least 70% of their values available, or less than 30% NaN values, had their NaN values filled in with summary statistics. The column median was used for this interpolation if the graphed distribution of a column showed that the distribution was heavily skewed or if it showed that the distribution was bimodal or multimodal. The mean was used to fill this subset of NaN values otherwise. A small exception was made for a set of 3 Seagate BarraCuda SSD model drives that had 8 columns with NaN values where no other models did. For these, the summary statistics to fill in their NaN values were calculated from the subset of rows that matched the drive models' capacity and manufacturer rather than the column total. This approach provided simple but more specialized approach to filling the missing values for this subset that had a minimal amount missing.

The remaining SMART attribute columns were over the threshold of missing values, so categorical variables were created to preserve the information that was available. The same rules as before were applied to determine whether the mean or median was used to determine class. Categorical column class was assigned based on if the row value was NaN, below the column summary statistic, or above the column summary statistic. Finally, column correlation was checked, and a few columns including model were dropped to avoid collinearity issues in the predictive modeling.

20 continuous and 11 categorical SMART attribute columns remained along with the capacity and manufacturer columns as the set of study factors for analysis. Pearson correlation coefficients for the continuous variables and Fisher's exact tests with Monte Carlo simulation for

the categorical variables were calculated. Once these were performed, the categorical columns were split into encoding columns and the data was split into 70% training, 20% testing, and 10% validation datasets via stratified sampling to preserve the failure class ratio. The continuous variables were standardized with scikit-learn's StandardScaler to ensure proper model training. Once the scaler was fitted to the training data, it was used to standardize the testing and validation data.

PCA was used to transform the continuous variables into their principal components, and a scree plot was used along with a cumulative inertia plot to determine that 13 of the 20 components, 65% of them, made up over 82% of the training dataset inertia. The other 7 components were dropped from the dataset to speed model training and testing, and to simplify the models while losing as little information as possible. Once fitted on the training data, the PCA was then used to transform the testing and validation data. Finally, SMOTE was used to oversample the training failure instances to a 1:1 ratio so that predictive models could learn to predict failure and not only non-failure.

Given the study factors significance, 6 predictive models were trained to predict HDD failure: An LBFGS solved logistic regression, a best-split decision tree with a maximum depth of 20, a non-weighted random forest ensemble with a max depth of 20, a class-weighted random forest with 1:2 majority:minority weighting with a max depth of 20, and 2 fully connected feed-forward DNNs, one with a simple architecture and another with a more complex architecture. The DNNs used the Binary Cross Entropy Loss, or BCELoss criterion, the Adam algorithm with a learning rate of 1e-7 as the optimizer, and the weights initialized with Xavier, or Glorot, initialization.

The Project Findings

The analysis showed that all of the study factors are statistically significant and useful for predicting HDD failure before it occurs. Based on the results of the exploratory data analysis and the predictive analysis, the SMART attributes 5, 197, and 9 are the best predictors for HDD failure. The manufacturer is also very influential when combined with the other data. For predictive models, the logistic regression model or the more complex MLP neural network are the two best model approaches for attempting to automate a data center's approach to predicting HDD failure before it happens.

Table 1

HDD Failure Predictive Model Testing Results

Model	Sensitivity	Specificity	Precision	Error Rate	ROC AUC
Logistic Regression	**0.6397**	0.9732	1.1478e-3	2.68%	0.8729
Decision Tree	**0.4412**	0.9690	0.8829e-3	3.10%	0.6900
Random Forest	**0.3603**	0.9903	2.2900e-3	0.98%	0.7974
Class-Weighted Random Forest	**0.4044**	0.9717	0.8858e-3	2.83%	0.7998
Simple DNN	**0.6176**	0.9185	0.4696e-3	8.15%	0.7681
Complex DNN	**0.7132**	0.9364	0.6946e-3	6.36%	0.8248

The Project Limitations

A few limitations of this project exist. First, a very large amount of the dataset was made up of missing values. A second limitation that deserves caution is that the ratios of drives made by each manufacturer in the dataset is very imbalanced. No assumptions about value or reliability of the four manufacturers included in the dataset should be made from this data. A third limitation is that the dataset was extremely imbalanced in terms of the minority (failure) and majority (non-failure) classes. Though SMOTE succeeded exceptionally well at allowing predictive models to learn from the imbalanced data, it does introduce bias as the synthetically created instances of the minority classes overrepresent their information in the analysis. Finally, working computer memory was a great limitation throughout the project as the dataset is so large. This limitation prevented factor analysis of mixed data from being performed and PCA had to be selected as the alternative.

Actions Proposed and Expected Benefits

It is highly recommended that either the logistic regression model or the more complex DNN model is added to the daily HDD diagnostics checks and backup procedure pipeline. The complex DNN will successfully flag 71.3% of drives expected to fail that day and the logistic regression 64%, allowing for total backup and retirement of the drive before the failure occurs. Do note that while more sensitive to detecting failure, the DNN does have a higher false positive rate, at 6.36%, than the more conservative logistic regression at 2.68%. Until this can be completed, special care should be given to drives with higher values of SMART attributes 5, 197, and 9 to reduce data loss and complications arising from the events of HDD failure.

Once implemented, an ensemble approach between the two should be tested to further reduce the false positive rate. Furthermore, additional research is warranted beyond the scope and limitations of the project. Taking an RNN approach to the data tidying and predictive modeling will almost certainly improve the results quite significantly, as they are specifically designed for time-series data such as this.

References

Backblaze. (2020). data_Q4_2019. San Mateo, CA; Backblaze.

Painchaud, A. (2018, October 31). *8 Reasons on How Data Loss Can Negatively Impact Your Business*. https://www.sherweb.com/blog/security/statistics-on-data-loss/.

Weiss, G. M. (2013). Foundations of Imbalanced Learning. *Imbalanced Learning*, 13–41. https://doi.org/10.1002/9781118646106.ch2