

A practical look at predicting Remaining Useful Life of Hard Disk Drives using LSTM Networks on SMART data:

Literature review and critique of findings

Cristian Seceleanu
Student Nr.: 29802635
Department of Computer Science
University of Reading
Reading, United Kingdom
c.seceleanu@student.reading.ac.uk

Abstract—As the demand for physical and cloud storage is increasing rapidly, the number of hard disk drives (HDDs) in operation is also increasing and with it so is the number of disk failures. As these failures usually impact the quality of the storage services it is clear that these cannot be ignored and a more proactive approach is required (waiting for a disk to fail before replacing it is more disruptive than replacing the disk before it is about to fail). This paper looks at several state-of-the-art approaches for predicting the remaining useful life (RUL) of HDDs using Long Short-Term Memory (LSTM) networks on Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) data.

Keywords—*Predicting Remaining Useful Life, RUL, Long Short-Term Memory, LSTM, Hard Disk Drives, Machine Learning, disk failure*

I INTRODUCTION

Hard Disk Drives (or HDDs) were introduced by IBM in 1956 and since then they have become the most wide-spread technology for data storage. They remain the most popular storage media in data. They remain the most popular storage media in data centers even after the rise of the Solid State Drive (or SSD) – which no longer has moving parts but rather chips with storage cells – because of their price to capacity and life expectancy ratio [2].

With the wide spread adoption of cloud services for workloads ranging from small (ie. individual virtual machines hosting a personal blog) to big (ie. using data science to predict or model weather patterns based on massive data sets collected over decades) it becomes clear that storage systems are required to scale to Petabytes and Exabytes which results in using hundreds of thousands and millions of HDDs per data center. At this scale disk failures are no longer rare events but rather they become the norm and with that comes the need to have optimal strategies to deal with such failures.

It is true that data loss caused by disk failure has been reduced by the adoption of solutions such as redundant arrays of inexpensive disks (RAID) however, when a disk that is part of a storage array fails and is replaced, the recovery process is a lengthy one and while it is running, additional stress is added on the remaining disks which can cause, in the best case scenario, performance degradation of the system,

and, in the worst case scenario, data loss caused by the failure of one or more disks in the same storage array. This approach works however due to its reactive nature it remains an unsatisfying solution [2].

In recent years focus has been shifted towards exploring more proactive solutions such as predicting when a HDD is close to failure such that the maintenance window required to replace it can be scheduled in advance to reduce the impact on the overall performance of the system [2, 5].

Due to shifts towards predictive systems, machine learning approaches have been gaining increasing popularity – especially the ones using models trained on S.M.A.R.T. data by relying on internal attributes of HDDs as indicators of drive reliability [2].

The objective and scope in this study, as the results of literature selection process mentioned in the next section, is to evaluate several state-of-the-art approaches to predicting RUL of HDDs using the LSTM neural network, that use the Backblaze data set, for practical application.

II LITERATURE SELECTION

The selection of the literatures was a lengthy one and started by searching for articles on predicting HDD failure using S.M.A.R.T. data and machine learning on Google to get a first feel for how popular this particular topic is (this produced a search result that was hundreds of pages long). While reading some of the materials found this way a second issue was identified and that was replicability of results which would not be possible without having access to a public data set (ie Backblaze [1]) to test the findings on. The next step was to search for articles on IEEE and Science Direct on the same topic only this time the purpose was to identify state-of-the-art approaches. Given that the search was a bit generic and that this has been a hot topic for years the result list was again very long (hundreds of articles) covering many approaches from using Decision Trees and SVM to predict the state of a HDD to predicting RUL using neural networks. For this paper the most recent work in the field that reported accuracy results close to or over 90% when predicting RUL was chosen which reduced the list to less than 10 articles published in the last 2 – 3 years out of

which 4 were selected based on writer experience in the field and relevance in context of predicting RUL for HDDs.

The initial understanding of the problem was that it would be a classification problem (ie. when certain S.M.A.R.T. attributes reach certain values it means the HDD will fail “shortly” – without actually knowing what “shortly” means – hours, days, weeks) however after reading the work in the field it becomes clear that in order to predict when a disk will fail one needs to look at the history of the HDD and make the prediction based on that (to transform “shortly” to an actual number of hours, days, weeks) hence the selection of LSTM network which looks at the evolution of the recorded data over time to make the prediction.

III LITERATURE REVIEW

The literature review is mainly based on four literatures. Hu et al. [5] proposes a model based on LSTM to predict disk failure in a given interval (30 days before the actual failure). Santo et al. [2] follows recent research in predictive maintenance, provides an overview of State-of-the-Art approaches and presents a deep learning approach to address data sparsity, need for domain knowledge and feature engineering to predict RUL of a HDD by identifying specific health conditions on the basis of S.M.A.R.T. attributes values using three main steps: defining the health degree for each HDD, extracting sequences in a specific time window for each hard disk and then assessing the health status through LSTM by associating a health level to each temporal sequence. The Conf. Paper [3] proposes a fault prediction method based on multi-instance LSTM neural network where the data in the entire degradation process is regarded as a sample then using the LSTM network the time characteristics of the data are mined and finally a multi-instance learning method is used to treat the degradation characteristics of the full-life data as a data bag and divide it into multiple instances thus the entire life cycle data is used for HDD abnormality detection. Coursey et al. [4] proposes methods for data standardization, normalization and RUL prediction using Bidirectional LSTM network with multiple days of look-back period considering S.M.A.R.T. attributes highly correlated to failure and builds a prediction pipeline that takes into consideration the long-term temporal relations in the failure data.

III.A Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.)

S.M.A.R.T. is a self-monitoring system supported by disk manufacturers in their products. The system detects and reports certain indicators correlated with disk failure with the detection algorithm being threshold based triggering a warning when any S.M.A.R.T. value exceeds its predefined threshold (which is set by the manufacturer). Previous measurements have shown that its fault detection rate is only 3% - 10% with a false alarm rate of about 0.1%. To enhance the performance of disk failure prediction based on S.M.A.R.T. machine learning and statistical techniques are proposed for building the prediction models with some of the most recent being based on LSTM networks which take into account the temporal evolution of data.

III.B Long Short-Term Memory (LSTM)

Long Short-Term Memory (LSTM) is an artificial neural network capable of capturing dynamic temporal behavior in time series by using shared parameters while traversing through time. The common LSTM unit is composed of a

memory cell, an input gate, an output gate and a forget gate. The input gates control the flow of input activations into the memory cell. The output gates control the output flow of cell activations into the rest of the network. The forget gates address a weakness of LSTM models which would otherwise prevent them from processing continuous input streams that are not segmented into subsequences. The forget gate decides which information needs to be memorized or not by taking in information from the previous cell and the current input. Whatever information is kept goes through the input gate. This determines what values will be updated in the cell. The \tanh function is applied on the cell state and current input for regulation. The cell state is then updated according to the combination of forget and input gates. Using the current cell gates and state, the output gate decides what to pass on to the next cell. A diagram of the LSTM cell is shown in Fig. 1.

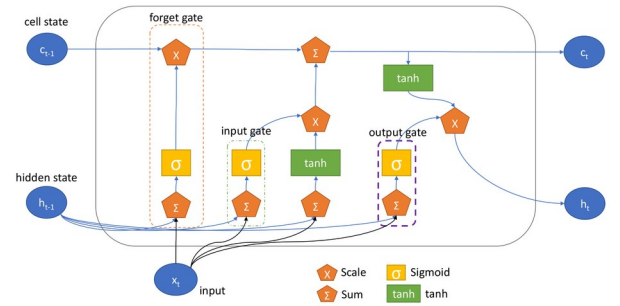


Fig. 1. The Structure of an LSTM Cell [4]

III.C Bidirectional LSTM (Bi-LSTM)

A bidirectional LSTM is a variant of an LSTM that consists of two LSTMs which run at the same time. One runs forward on the input sequence and the other runs backwards on the input sequence. In the context of this paper one could think of one direction being LSTM running on the sequence of HDD data leading up to failure and the other direction being running backwards (away) from disk failure. This allows the LSTM to better learn the relationship between the attributes and the RUL with a simple architecture change. The architecture of a Bi-LSTM is shown below in Fig. 2.




Fig. 2. Bidirectional LSTM Architecture [4]

III.D Data Preprocessing

The approach presented in Santo et al. [2] for data preprocessing consists of Feature Selection and Health Degree Computation where each data set is treated separately. For the Baidu data set, where S.M.A.R.T. data is available for 20 days before disk failure they propose a model that will predict the health status 20 days in advance (depending on the splits of the selected Regression Tree

model of the feature Time-to-failure) and for the Backblaze [1] data set, where S.M.A.R.T. data is collected over a long period of time, several lengths for the prediction window are tested (15,30,45). For the techniques used 70% of the data is used for training, 15% for testing and 15% for validation.

In the Conf. Paper “A multi-instance LSTM network for failure detection of hard disk drives” [3] the data sets contain S.M.A.R.T attributes from a communications company (SAS HDDs) and from Backblaze [1] (SATA HDDs). Because SAS and SATA HDDs have different S.M.A.R.T. attributes (SCSI vs ATA standards) only 6 were selected. Sixty percent of the life cycle of the HDD is used as a positive sample with normal and fault samples coming from different hard drives.

Coursey et al. [4] selected only one HDD model (ST4000DM000) from the Backblaze [1] data and programmatically created a data set containing the S.M.A.R.T. features leading up to failure. For feature selection two methods were used: a correlation score and a Decision Tree but before using any of these all null values features as well as normalized features were removed which led to 5 S.M.A.R.T. attributes being selected.

Hu et al. [5] selected only two types of disks (ST4000DM000 and ST8000DM002) for their experiment because the S.M.A.R.T. data collected for these models is relatively large in the Backblaze [1] data set. In the disk failure prediction problem they refer to the failed disk as a positive sample (with the healthy disk being a negative sample) and keep all positive samples in the data set and use down sampling for the negative samples so that the ratio of negative to positive samples is maintained at about 1 to 4. Before further analysis only raw attributes are used then use the Pearson Correlation on each feature to indicate if it can differentiate between the positive and negative samples thus reducing the number of attributes to 10 for the training of the model. To avoid bias towards features with larger values zero-mean score is applied for data normalization.

III.E Results

For Santo et al. [2] when compared their best result on the two data sets with other state-of-the-art methods outperforms all models in terms of accuracy on failed sequences, FDR (Failure Detection Rate) and FAR (False Alarm Rate) for both HDD health status assessment and HDD failure prediction.

The multi instance LSTM model presented in Conf. Paper [3] achieves state-of-the-art performance with low FAR even on long time series.

The Bi-LSTM model from Coursey et al. [4] achieves high accuracy (state-of-the-art) within the same range it is trained on and also outperforms LSTM when predicting RUL

In Hu et al. [5] the proposed LSTM architecture performs well in sequential disk failure prediction with a reported precision of more than 85% and only 1.3% FAR.

IV FINDINGS

Having reviewed a number of papers it is clear that using LSTM networks for predicting RUL for HDDs has the highest accuracy however in order for such models to have a practical application, which is the scope of this paper, there are a few aspects that need to be addressed:

IV.A The hardware manufacturing process

Each HDD model is unique due to the fact that its production line is different from model to model, the materials used for making it are also different (some HDDs have Helium inside rather than air) as well as other factors about it (for example handling, packaging, etc) which means under the same workload two different HDD models with similar specifications will have different wear levels and thus will fail differently which in turn means the training model needs to take this into account and given that manufacturers do not share specifics about how various HDD models are built the models need to be trained with data collected from each and every disk type we want to predict RUL for.

IV.B The workload

The literature selected for this paper uses the popular Backblaze [1] data set to train the models (some train the models on extra data from Baidu and other companies) however, the fact that there is no information in these data sets about the workload itself, accuracy will differ when training a model on the Backblaze [1] data set (where the workload is generally backup) then use it to predict failure or RUL on a data set collected from say a cloud service provider from virtual machine hosts (which might be used for heavy local disk operations or have a very low disk IO operation workload) or on a data set collected from a University from its computer lab where students have courses a number of hours per week after which most if not all machines are powered off or suspended.

IV.C Keeping the model up to date

New models of HDDs are made available to consumers every year and with that comes the need to retrain the models to take these into account. The approaches seen so far in the literature that was selected for this paper cover training the models however there is also a need to keep the models up to date by constantly retraining or by using continuous learning such that when a new disk is introduced into a system the model can get real-time training. With the ever increasing demand for storage capacity and performance new disk models are being launched every year and service providers such as Backblaze start adopting them to expand their storage capacity and increase performance for the services that they provide however in doing so the prediction model also needs to be updated such that it can predict failure for the new disk models.

IV.D Cannot rely only on historical data

Any practical application of such a prediction model is incomplete in the sense that if it only relies on historical data belonging to a third party entity it cannot be applied with the same level of accuracy to a new entity unless the new entity uses the exact same disk model and workload. In practice a model that uses a hybrid approach such as classifying the disk health status based on current S.M.A.R.T. attribute measurements at first (ie. good, fairly good, bad etc) then as more and more data is collected from the disks it transitions to the RUL prediction model based on historical information would be more appropriate.

V REFLECTION

The study objective of this paper has been defined with appropriate level of specificity.

Approaching the literature selection process as described in this paper provided a very good starting point for

understanding the problem and identifying relevant works in the selected field, the only difficulty being the sheer volume of information that needed to be screened and evaluated as relevant for this paper's context. This systematic approach worked so well that, if asked to re-write this paper or write new articles in the future, it would still remain the preferred one.

The literature selected for this paper is state-of-the-art and can be found easily. The authors of the papers have participated in writing other storage related works (including predicting disk failure) and their works have been cited in various other papers.

The four articles were reviewed in detail and are relevant for the selected topic, however there are considerably more papers that have been reviewed and mentioned many times in other works before, which can be used for practical applications in classifying the health state of HDDs and, based on that, predicting if a failure is imminent.

REFERENCES

- 1 A. Klein, "Backblaze Drive Stats for Q2 2022", Aug. 2019, [online] Available: <https://www.backblaze.com/blog/backblaze-drive-stats-for-q2-2022/>.
- 2 A. De Santo, A. Galli, M. Gravina, V. Moscato and G. Sperli, "Deep Learning for HDD Health Assessment: An Application Based on LSTM," in *IEEE Transactions on Computers*, vol. 71, no. 1, pp. 69-80, 1 Jan. 2022, doi: 10.1109/TC.2020.3042053.
- 3 "A multi-instance LSTM network for failure detection of hard disk drives," *2020 IEEE 18th International Conference on Industrial Informatics (INDIN)*, 2020, pp. 709-712, doi:10.1109/INDIN45582.2020.9442240.
- 4 A. Coursey, G. Nath, S. Prabhu and S. Sengupta, "Remaining Useful Life Estimation of Hard Disk Drives using Bidirectional LSTM Networks," *2021 IEEE International Conference on Big Data (Big Data)*, 2021, pp. 4832-4841, doi: 10.1109/BigData52589.2021.9671605.
- 5 Lihan Hu, Lixin Han, Zhenyuan Xu, Tianming Jiang and Huijun Qi, "A disk failure prediction method based on LSTM network due to its individual specificity", *Procedia Computer Science*, vol. 176, pp. 791-799, 2020.