21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2018 Grid of the Future Symposium

Latent Variable Survival Models for Network Transformer Prognostics

J. GARRITY[1], B. KITTRELL[2], S. MCCORMICK[1], W. VEGA-BROWN[1] *
Tagup Inc.[1], Con Edison[2]
USA

SUMMARY

We introduce a statistical approach to incorporating high dimensional, high resolution remote monitoring system measurements into a predictive model of the remaining useful life of transformers. Popular tools for survival analysis, such as the Kaplan-Meier estimator and the Cox proportional hazards model, cannot easily incorporate time-series data or model degradation over time. In contrast, our model computes a statistical summary of all historical measurements and uses that summary to predict the risk of future failure, enabling robust estimates of remaining useful life and probability of failure in a given year. We scale to very large data sets by using a combination of feature extraction, data compression, convex optimization, and parametric models. We demonstrate the utility of our approach using historical operating data from the Con Edison network transformer fleet. Our case study covers the approximately three thousand network transformers on the island of Manhattan, many of which have instrumentation measuring load, temperatures, and oil level, among other variables. We demonstrate our approach using years of network transformer sensor data across thousands of network transformers to build a model capable of predicting the remaining useful life of each network transformer in the Con Edison fleet.

KEYWORDS

Network transformers, asset management, survival modelling, latent hazard model, machine learning, failure prediction

* We are grateful to the New York State Energy Research & Development Agency (NY-SERDA) for supporting this work through PON 3397, without which this study would not be possible.

jon@tagup.io

Introduction

Con Edison maintains a fleet of approximately 27,000 network transformers to distribute power throughout the city of New York. Though transformers are generally highly reliable, across a fleet of this size unexpected (and sometimes catastrophic) failures are almost certain to occur each year. These failures[1] are expensive, both in terms of downtime, manpower, and equipment cost—and some of that cost could be avoided with a better understanding of when a transformer was likely to fail.

Better understanding when or how an asset will fail creates business value across many asset classes, network transformers included. First, preemptive maintenance or removal can reduce in-service failures, avoiding unplanned downtime, premature asset loss, negative safety impact, and expensive collateral damage. Second, accurate estimates of remaining useful life prevents premature replacement of an asset, reducing amortized capital expenses. Third, estimating failure probabilities allows for optimized field force strategies, providing a higher “reliability return” for each inspection. Finally, understanding what assets will fail in an upcoming budget cycle ensures that there is adequate lead time for replacement, enhancing inventory management and procurement.

Existing strategies for preventing in-service failures rely heavily on operator experience and or on ad hoc heuristic estimates. For example, a common strategy relies on subject matter experts to review operating data, inspection reports, and asset metadata, such as the transformer manufacturer and install date. This is labor-intensive; with tens of thousands of transformers, the expert may not have time to analyze each transformer in detail. To improve productivity, operators may rely on threshold alarms to prioritize assets. If, say, dissolved gas levels or rates of change exceed a threshold, a transformer may be flagged for attention and preemptively removed from service. Defining these thresholds is challenging: if the thresholds are set too high, critical issues will go unexamined. If they are set too low, experts will be unable to investigate each alarm.

To make informed decisions to inspect, repair, or replace an asset, experts rely on heuristics derived from their experience. New methods try to capture some of this expert knowledge by coding asset health indices (AHI). Typically, these indices are linear combinations of observables, often scaled as a percentage. The index, while useful for rank ordering assets for replacement or inspection, are difficult to interpret: what does it mean to be 85% healthy? Furthermore, there is no established AHI standard or consistency in definition, and an unclear relation between an AHI and failure probability. We overcome these limitations by leveraging a recent innovation in survival analysis, latent variable hazard modeling, to model both phenomena that lead to an instantaneous increase in the risk of failure, such as overloading, and phenomena that degrade the long-term health of the asset, such as corrosion.[2] By using statistical tools to derive measures of asset health, we can make principled, interpretable, and verifiable estimates of useful quantities like remaining useful life or probability of survival until some future time. Moreover, unlike conventional health indices, our estimates improve as we collect more training data, allowing us to discover connections and patterns beyond what the human eye can discern. In the next section, we provide an overview of survival analysis and describe our model in detail; in the following section, we describe a case study with the Con Edison network transformer fleet.

[1] Failures here are restricted in definition to in-service failures (as opposed to pre-emptive removals).
[2] See Moghaddass & Rudin (2015) for a more thorough review.

2. Survival Analysis

2.1 Overview

We wish to predict the risk of failure across a fleet of transformers. If transformer n was installed at date $t_{i,n}$ and failed at date $t_{f,n}$, we say it has a lifetime $T_n = t_{f,n} - t_{i,n}$. The objective of survival analysis is to understand the distribution over the lifetime of a transformer, given static metadata (transformer rating, manufacturer, etc.) as well as operational data (such as temperature and loading). In the remainder of this section, we first introduce a useful statistical tool—the hazard function—which we will use to model the distribution over possible lifetimes of a given asset. We then describe the limitations of conventional approaches for estimating the hazard function, and finally we will describe the latent variable model for estimating the hazard function, which overcomes many of these limitations.

2.2 The hazard function

We may be interested in a variety of statistics related to the lifetime T_n of transformer n. For example, we may wish to know the expected remaining useful life $m_n(t) = E[T_n|T_n > t]$; this is the average age at which the transformer is expected to fail, given that it has been in service for t days. This quantity generalizes the mean time between failures by accounting for the operational history of the transformer. We can also estimate the uncertainty of the expected lifetime.

Alternatively, we may wish to understand the probability $(t_0 < T_n < t_1|T_n > t)$ that a given transformer n fails in a specific time window (i.e. between t_0 and t_1), given that it has been in service for t days. We could use this to prioritize maintenance on those transformers most likely to fail within the next month. Additionally, this probability could be used for purchasing decisions by estimating the fraction of transformers that will fail in the next year. In addition, from this probability, we can derive statistics like the *survival function*, $S(t,\ t') = \Pr(T_n > t'|T_n > t)$, the probability that a transformer remains in service for t' days given that it has survived for t days.

It turns out that these and other statistics can be captured by a function called the hazard function[3] $\lambda_n(t)$, which is the probability of a transformer failing during the t^{th} day of its life, given that it remained in service until day t:

$$\lambda_n(t) = \Pr(T_n = t + 1|T_n > t)$$

We can derive the expected lifetime from the hazard function:

$$m_n(t) = \mathrm{E}[T_n|T_n > t] = \sum_{u=0}^{\infty} \exp\left(-\sum_{\tau=t}^{t+u} \lambda_{n,\tau}\right)$$

We can also derive the survival function or failure probability in a given period:

[3] Formally, the hazard function is a continuous function of time $\lambda_n(t)$ representing the density of failures at time t; the probability of failure within a given time period from t_0 to t_1 is given by integrating the hazard function from t_0 to t_1. We use a discrete model for didactic clarity.

$$S_n(t) = \exp(-\sum_{\tau=0}^{t} \lambda_{n,\tau})$$

$$\Pr(\tau_0 \leq T_n < \tau_1 | T_n \geq t) = \exp(-\sum_{u=t}^{\tau_0} \lambda_{n,u}) - \exp(-\sum_{u=t}^{\tau_1} \lambda_{n,u})$$

In fact, the hazard function $\lambda_n(t)$ is sufficient to calculate any statistic derived from the lifetime of an asset. Unlike derived statistics, however, the hazard function allows us to use *both positive information* (the knowledge that a transformer failed at age T_n) and *negative information* (the knowledge that it did *not* fail before time T_n) independently. This in turn makes it easier to tease out the impact of a particular time series measurement on the hazard, and thus on the remaining useful life. In the following sections, we will describe several conventional ways of estimating the hazard function from data and discuss the limitations of each.

There are many ways to model the hazard function. The simplest is to assume it is constant, leading to an exponential survival function:

$$\lambda_n(t) = a$$

$$S_n(t) = \exp(-at)$$

There are a variety of parametric models that improve upon the constant-hazard model, such as the Weibull or log-normal models; the Weibull model assumes the hazard is a polynomial:

$$\lambda(t; a_n, b_n) = \frac{a}{b}\left(\frac{t}{b}\right)^{a-1}$$

Which gives an implied survival distribution of:

$$S(t; a, b) = \exp\left(-\left(\frac{t}{b}\right)^{a}\right)$$

Although these simple parametric models are capable of estimating the distribution over lifetimes for a population of transformers, they cannot account for the idiosyncrasies of an individual transformer, nor can they easily account for differences in the operating histories of individual transformers. Often we have access to information like a transformer's manufacturer or rating, and this information should influence its expected lifetime.

Many authors have explored using time-series data for failure analysis. For example, the proportional hazards model developed by Cox (1972) was extended to incorporate time-dependent information by Fisher & Lin (1999). This class of models is effective for estimating the hazard function and the impact of observations on short term failure risk, but cannot model the long-term impact of equipment degradation, as it does not incorporate the impact of *past* observations on the *future* hazard estimates. Suppose we observe a transformer overloaded at time t and the model predicts the transformer is more likely to fail when overloaded. If the transformer survives the overloading, the Cox model assumes that overloading has no further impact on the expected lifetime. As the operational history grows longer, our understanding of failure risk should change; for example, an inspection report of corrosion should permanently increase the risk of failure and decrease the expected lifetime.

Models that explicitly include degradation often assume that the observed covariates have a direct relationship to degradation state, typically with additive random noise [e.g. Kharoufeh & Cox (2005), Zhou, Serban & Lin (2011)]. A third class of models assume that the degradation state is unobservable, but that the time-varying covariates indirectly reflect the degradation state over time. The degradation state can then be inferred using well-known approaches such as hidden or semi-hidden Markov models.

2.3 The latent variable model

The latent variable model introduced by Moghaddass and Rudin (2015) computes a summary of the impact of past observations. Additionally, the model estimates instantaneous and temporary changes to the risk of failure due to recent observations. Although the most likely parameters for this model cannot be found explicitly as they can in Cox-like models, we can find them efficiently by solving a convex optimization problem.

We assume that we have a body of historical records for many transformers, where the record for transformer n is a sequence $X_n = \{x_{n,t}\}$of measurements such as average load or temperature in a given period t. We also assume we have the observed lifetime T_n of the transformer.

The latent variable hazard model is defined by a vector of parameters θ and is given by:

$$\lambda_n(t; X_n, \theta) = \lambda_0(t)\left(\exp(\theta_1 \cdot x_{n,t}) + \sum_{\tau} \exp(\theta_2 \cdot x_{n,\tau})\right)$$

Here, $\lambda_0(t)$is a base hazard rate computed from the distribution of failure times across all units. This is scaled by a factor derived from the operational history of the transformer, as represented by the sequence $X_n = \{x_{n,t}\}$. In this work, we assume the base hazard rate is given by the Weibull model described in the previous section.

The relationship between the observed data and the hazard rate is encoded by the weights θ_1 and θ_2. The form of $\lambda(t)$ is chosen to allow us to model factors that represent both a short-term failure risk and a long-term health degradation. θ_1 models the relationship between the measurements on day t and the risk of failure on day t. For example, a transformer is more likely to fail under heavy load than when de-energized. In contrast, θ_2 models the relationship between failure risk on day t and all measurements prior to day t. For example, if we observe an unexpected drop in oil pressure or observe main tank corrosion on any day before day t, then the risk of failure on day t will increase.

Obtaining acceptable values for the parameters weights θ_1 and θ_2 by hand-tuning or by experiment would be difficult, as these values are not easily interpretable. Instead, we use optimization to find the most likely values given our collection of records $D = \{y_n, X_n, T_n\}$. The likelihood of a set of parameters is simply the probability of observing the measurements we recorded, and is given by:

$$\mathcal{L}(\lambda; \mathcal{D}) = \prod_n \left(\prod_{t=0}^{T-1} \exp(-\lambda(t; y_n, X_n)\right)(1 - \exp(\lambda(T_n; y_n, X_n))$$

Finding the parameters that maximize this function is a straightforward optimization problem[4] and can be solved using a variety of algorithms, such as gradient descent. We discuss the details of our method for solving this optimization problem in the next section.

2.4 Using incomplete data

[4] Specifically, the objective function is shown to be log-convex in Moghaddass and Rudin (2015).

In order to evaluate the hazard at time t, our model requires access to all observations up through time t. If t is a time in the future, this presents a problem. We do not know the future operating data. If we did, prognostics would be easy! In order to make predictions about the expected remaining useful life or the probability of survival until some point in the future, we must be able to predict the hazard without access to these future observations.

This problem is not limited to prognostication and prediction; even our records of the past operating history of each transformer are often incomplete. Many transformers have been in service since long before remote monitoring was available; sensors are retrofitted long after a transformer is put into service. Communication errors and electromechanical failures sometimes prevent data from being collected, and collected data may be lost or rendered inaccessible within an organization for a variety of reasons. In practice, a useful model must be able to make predictions using incomplete data.

We deal with this challenge by learning a fixed, constant risk with which we replace any missing data. This allows the model to approximate the expected hazard an observation if that observation is not available. Concretely, we replace any missing data with a fixed, constant value; this value can be set to the average of the available observations or can be estimated as part of the optimization problem. This allows us both to forecast into the future, and to address imperfect and incomplete historical data.

3. Case Study: Con Edison

3.1 Overview

Con Edison operates approximately 27,000 network transformers throughout the service territory. These network transformers are responsible for delivering reliable power to over three million customers throughout New York City and Westchester County. In an effort to reduce the number of in-service failures, a remote monitoring system was deployed to capture sensor data from a subset of the fleet for which the entire lifetime was observed. Data monitored includes loading, main tank pressure, temperature, and oil level. We restrict our analysis to an analytic sample of 243 transformers for which the installation and failure time of the asset was observed and for which remote monitoring system (RMS) data was available.

The project objective is to estimate the probability that each transformer will be removed from service before a given date. By estimating the hazard and survival functions of each asset, there is an opportunity to reduce in-service failures, enhance field force productivity, and augment procurement decisions.

3.2 Data

We test the applicability of the latent state hazard model using RMS data. The data consists of observations recorded at a frequency of approximately every ten minutes of variables such as transformer load, voltage, oil level, and temperature. In order to make estimating a survival model computationally tractable, we extract 127 summary statistics from the raw RMS data, calculated at a twelve-hour interval. The extracted

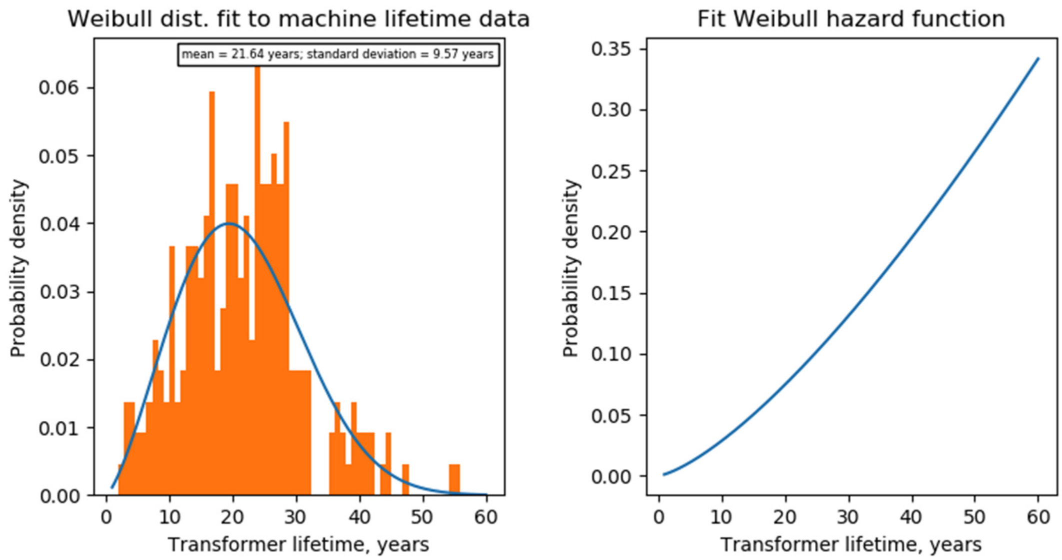

Figure 1. Estimating the base hazard rate from observed transformer lifetimes using the Weibull distribution (n = 243).

features capture information such as raw variable means, medians, and correlations within the twelve-hour interval.

3.3 Methods

To construct a latent state hazard model, we first estimate the *base hazard rate* λ_0 which captures average failure risk as a function of transformer age across the fleet. The base hazard rate is obtained by fitting the parametric Weibull model using maximum likelihood estimation, as shown in Figure 1. As the figure above shows, this model fit corresponds to a monotonically increasing hazard function, consistent with failure probability increasing with transformer age. The probability density on the vertical axis represents the predicted probability of transformer failure for any given lifetime between zero and 60 years.

3.4 Accounting for missing data

RMS data is often incomplete, due in part to the relatively recent introduction of the capability into the fleet. Therefore, many instances of missing data result when aggregating raw data into twelve-hour intervals via feature extraction. We first standardize the time series covariates by subtracting the mean and dividing by the standard deviation of each variable. Next, we fill all missing data with zero values, which corresponds to the mean in our newly standardized data. Filling missing values with the mean is roughly equivalent to assuming steady state machine operation. More sophisticated methods for dealing with missing values are a high priority for future research.

3.5 Parameter optimization

To optimize the parameters of the latent variable model, we used the L-BFGS-B optimization algorithm developed by Zhu, Byrd, Lu, & Nocedal (1997) to maximize likelihood (as described in subsection 2.3). One benefit of the latent variable survival model is that the likelihood function can be computed as a

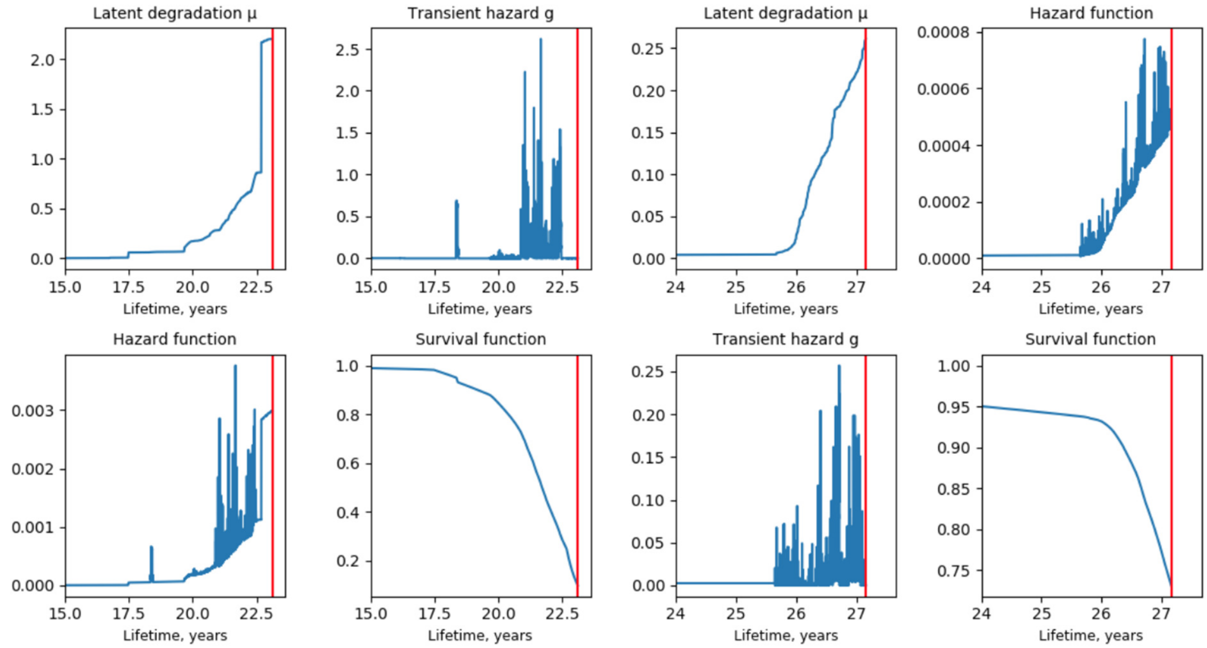

Figure 2. Components μ (latent degradation) and g (transient hazard) of estimated hazard function and corresponding survival function for two transformers in the validation set.

function of the observed data, meaning that standard optimization techniques can be applied to find the values of θ_1 and θ_2. To prevent overfitting, we also include L2 regularization [Hoerl & Kennard (2005)] in our optimization process. Regularization provides a way to explicitly penalize complex models and improves generalization performance on unseen data.

3.6 Validation and evaluation criteria

While maximum likelihood is a convenient method to determine the optimal values of the parameters of our model, our primary goal is to achieve high accuracy when predicting remaining useful life and failure probabilities. We selected metrics based on considerations of what would be most practically useful for prioritizing maintenance and minimizing in-service failure rates. After fitting our model on observed *training data*, we calculate these metrics on held-out *validation data* to get an unbiased measure of model performance.

As illustrated in subsection 3.7, our model is able to identify subtle signals indicating proximity to failure, using only RMS data up until the current period. This enables real-time renderings of the hazard and survival function. In the following section, we present such a method to predict failure at a one- and two-year horizon. Our evaluation of model accuracy is summarized through *precision*, the percent of predicted failures that actually resulted in failures, and *recall*, the percent of actual failures that were correctly predicted. The model allows the user to tune the failure prediction framework to their specifications, based on their relative aversion to *false positives* (when failure is predicted by the model, but no failure occurs) and *false negatives* (a failure occurs, but was not predicted by the model).

3.7 Results

After training our model on 125 historical transformer records, we validated its performance on 20 test transformers. By plotting the hazard and survival functions for several of the test transformers, we can see

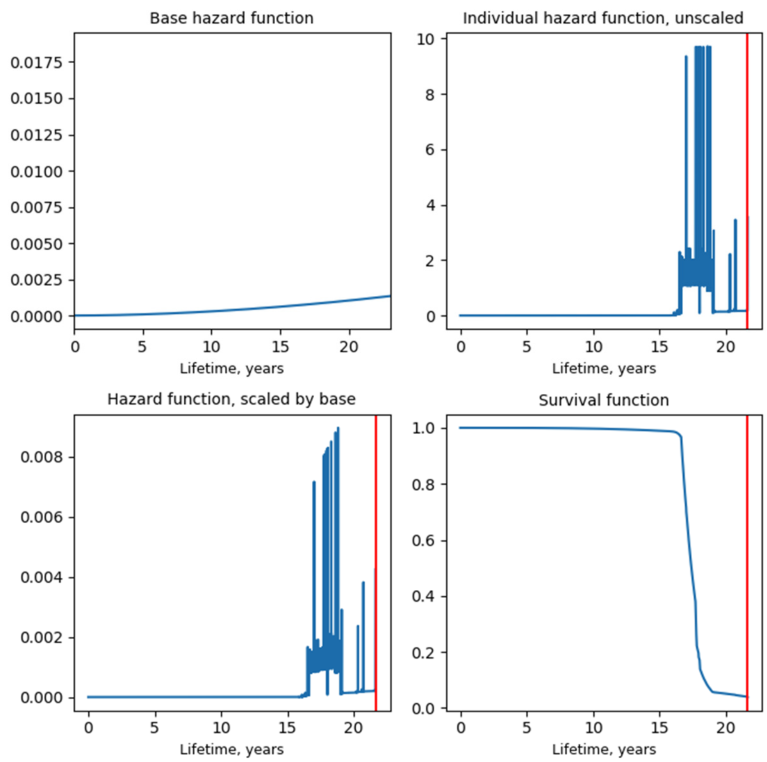

Figure 3. An illustration of estimating the survival function by multiplying base hazard and individual base hazard for an example transformer in the validation set.

clear patterns in the data indicating an upcoming failure. These are reflected intuitively as a sharp drop in the corresponding survival function. More concretely, our failure prediction classification model achieved a precision of 0.73 and recall of 0.40 when forecasting transformer failures at a time horizon of one year, and a precision of 0.57 and a recall of 0.65 at a two-year horizon.

As described in section 2, the latent variable survival model consists of both a degradation term μ and a transient hazard term g. The sum of these components, together with the base hazard, provides the hazard function (probability of failure at each time t, given that failure has not yet occurred) from which we derive the survival function (probability of the transformer remaining in service given that it has lived until time t). Figure 2 illustrates the components of the hazard function and how they inform both the hazard and the survival function. The plots of μ, g, the hazard function, and the survival function are shown for two transformers in the validation set. For both of these transformers, there is a significant disturbance in the hazard function during the last two years of their operation, suggesting an imminent in-service failure.

Once the transformer-specific hazard function is constructed from the latent degradation term and the transient hazard term, it is multiplied by the base hazard rate as described in the Methods section (see Figure 1) This ensures the hazard function increases over time, even during periods with missing RMS data. Figure 3 shows this process for a transformer in the validation set. This transformer showed signs of increasing failure potential almost five years before it failed in-service.

For an engineer or maintenance worker, it would be beneficial to understand the underlying causes of failure. To this end, we identified features that took on unusual values (at least two standard deviations from the mean) during and just before the period when the hazard function and the degradation function μ first

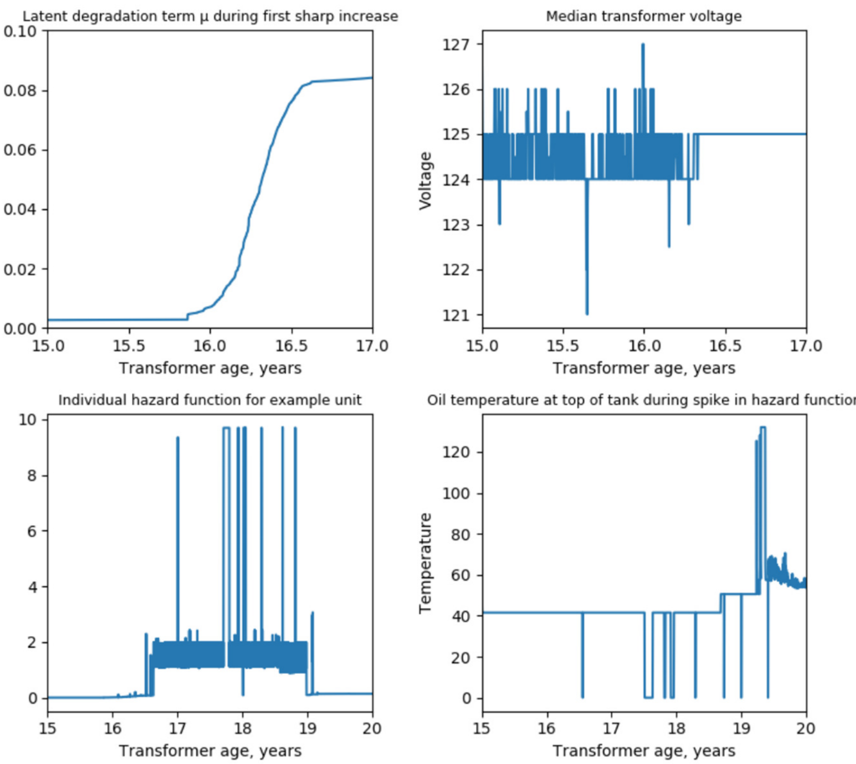

Figure 4. Plots of time series covariates with unusually large deviations from the mean during large spikes in the hazard function (bottom left) and the latent degradation term (top left) for an example validation transformer. Unusual fluctuations of voltage (top right) accompany the first large increase in the degradation component μ. In contrast, the spikes in the overall hazard function are accompanied by irregular values of transformer temperature measured at the top of the tank (bottom right). It appears that unusual covariate values are manifested first through the degradation term μ before being accounted for by the transient hazard term g and the overall hazard function.

increased dramatically. The top panel of Figure 4 shows the individual hazard function for a validation transformer along with several of the variables that were unusually high during the period when the individual hazard function began to spike. The bottom panel shows an analogous plot for the degradation function associated with the same transformer, which begins to increase before any significant change in the hazard function occurs.

The signs of increasing failure probability are manifest in the hazard and survival functions. We used our training data to compute estimated failure probabilities at both the one- and two-year horizon, compared it to the ground truth (i.e., whether the transformer was in fact removed before the query date) and chose the classification threshold that minimized the number of incorrectly classified transformers. We then computed one- and two-year failure probabilities on the transformers in the validation set and used the threshold from training to implement our classification rule as would have to be done in practice for machines that have not yet failed.

Figure 5 shows the receiver operating characteristic (ROC) curves of failure predictions at both a one- and two-year horizon. The curves illustrate the tradeoff between decreasing false positives at the expense of increasing false negatives. A perfect classifier achieves an AUC score (area under the ROC curve) of 1. Both our classifiers achieve scores of 0.92 on the validation data.

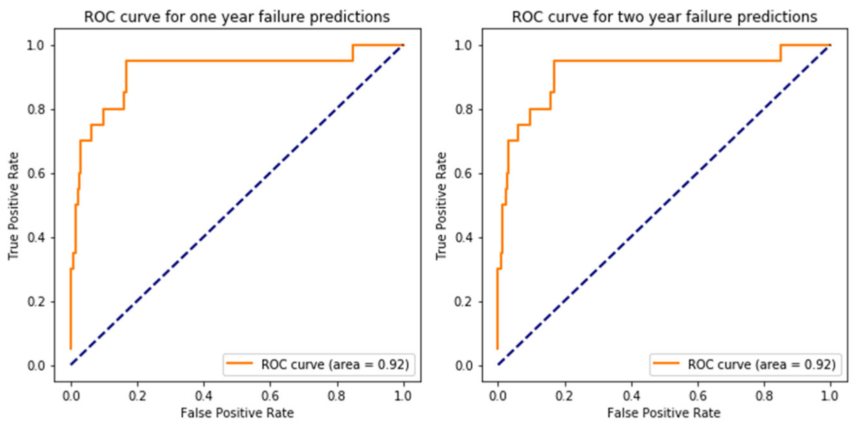

Figure 5. Receiver Operating Characteristic (ROC) curves for failure predictions at one- and two-year frequencies. The curve traces out the trade-off between the rate of false positives and true positives as the probability threshold for classification is increased.

4. Conclusions

We demonstrate the applicability of the latent variable survival model to 243 network transformers installed throughout Manhattan. This required first extracting time series features from the observed data, and accounting for missing data. We optimized the parameters of our model using 125 transformers and evaluated the performance on 20 validation transformers. Graphical representations of the hazard and survival function for many of the test transformers reveal periods of high transformer stress that anticipate failures before they occur. We used this preemptive signal of impending failure to develop a failure prediction model at both a one- and two-year time horizon, which achieved precision and recall scores of 0.73 and 0.4 (one-year horizon) and 0.57 and 0.65 (two-year horizon) respectively. Both the precision and recall of the model can be improved with more training data and enhanced feature engineering, which will be implemented in the next phase of the study. Additionally, incorporating transformer metadata would likely increase predictive accuracy.

With this model and implementation, we have created a general framework for assessing network transformer remaining useful life and failure probability. These methods are broadly applicable and can be used with other asset classes (e.g. circuit breakers, substation transformers). With achievable improvements in feature engineering and training data volume, this model will provide sufficient accuracy to optimize maintenance, procurement, and inspection strategies.

BIBLIOGRAPHY

Cox, D.R. (1972). Regression models and life-tables. *J. Roy. Statist. Soc. Ser. B*. **34** 187-220.

Fisher, L. D. and Lin, D. Y. (1999). Time-dependent covariates in the Cox proportional hazard regression model. *Annu. Rev. Public Health*. **20** 145-157.

Hoerl, A. E. and Kennard, R. W. (1970). Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*. **12** 55-67.

Kharoufeh, J. P. and Cox, S. M. (2005). Stochastic models for degradation-based reliability. *IIE Transactions (Insitute of Industrial Engineers)*. **37** 533-542.

Moghaddass, R. and Rudin, C. (2015). The latent state hazard model, with application to wind turbine reliability. *The Annals of Applied Statistics*. **9** 1823-1863.

Zhou, R. R., Serban, N. and Gebraeel, N. (2011). Degradation modeling applied to residual lifetime prediction using functional data analysis. *Ann. Appl. Stat*. **5** 1586-1610.

Zhu, C., Byrd, R. H., Lu, P. and Nocedal, J. (1997). L-BFGS-B-FORTRAN subroutines for large-scale bound constrained optimization. *ACM Trans. Math. Software*. **23** 550-560.