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

CIGRE US National Committee
2018 Grid of the Future Symposium

Machine Learning Techniques in Support of Better Asset Life Estimation - Let the Data Surprise You

L. CHEIM (1), D. DORRIS (2), S. SCHAAL (3), B. SINGH (4), T. GALOPPO (5)
(1, 3, 4, 5) ABB, (2) Nashville Electric
(1, 2, 4, 5) USA, (3) SWITZERLAND

SUMMARY

Machine Learning [1] is the part of Artificial Intelligence that utilizes a set of algorithms that are pre-programmed to learn from examples (i.e. a large number of power transformers) and make accurate predictions of future outcomes (i.e. a transformer is about to fail or has already failed) based on well trained datasets with multiple features (i.e. operational data of power transformers such as historical loading, dissolved gas analysis, oil quality, etc.). This paper shows the initial results of an exploratory data analysis as part of a research and development joint project between ABB and utilities in order to use historical information about healthy and failed power transformers and develop better estimates of life expectancy. The problem with existing life estimation models is the lack of real life correlation between theoretical approaches (i.e. Arrhenius ageing model in the case of power transformers) and failure of a given population of transformers. Besides, currently existing models arguably put a heavy weight on thermal aging alone whilst actual aging depends on a number of other factors, such as for example maintenance procedures and history, vintage, accessories, oil quality, cyclic loading and overloading practices, among others. Although at the initial stage of research the paper already shows some interesting and promising results obtained through the so called exploratory data analysis (EDA) that comprises a set of sophisticated visual statistical tools in preparation for the application of more complex predictive models.

KEYWORDS

Machine Learning, algorithms, asset management, life estimation

luiz.v.cheim@us.abb.com

1. Power Transformer Life Estimation

Figure 1 below illustrates typical estimated transformer life expectancy based on the solid insulation aging rate (Arrhenius model) and on the estimated effect of oxygen and moisture in addition to temperature as found in international documents and standards [2-4].

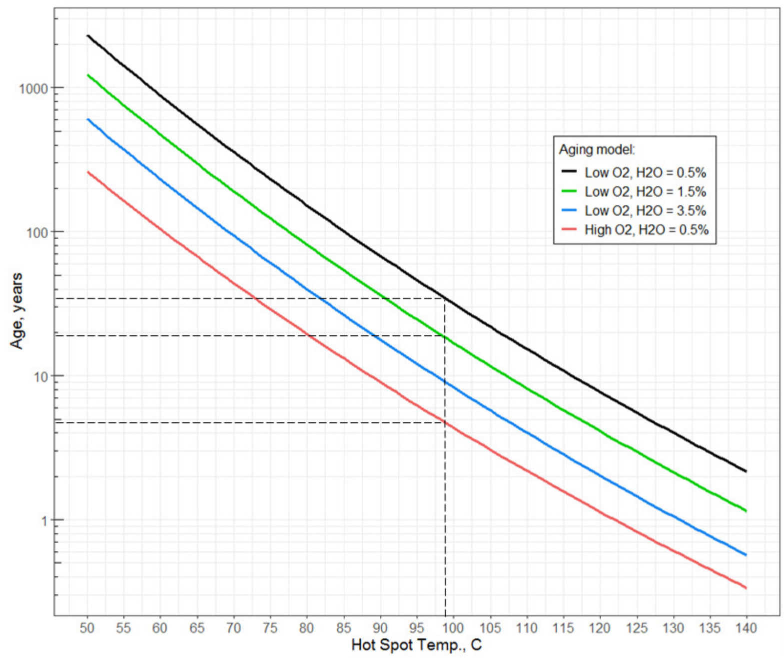

Figure 1 - Estimated effect of temperature, oxygen and moisture on power transformer aging (thermally upgraded paper) as a function of continuous hot-spot temperature, as given in [2].

The model shown in Fig. 1 is the best estimates that the industry has so far although some important limitations must be taken into account when the highly complex mechanisms of transformer insulation degradation (see Fig. 2) and practical life estimation are considered, as follows:

1. The model assumes that the age of the transformer is defined by the age of the hot-spot, a single point or a small region of one given winding where highest temperature is developed;
2. The model also assumes transformer operation at a continuous hot spot temperature, say 98C, as shown in Fig. 1 – this is almost never the case in real life;
3. The effect of oxygen and moisture are estimates based on laboratory investigations with the assumption that real life transformers will have the same or approximately the same effect;
4. Other important aspects of transformer design and construction (i.e. oil channels, oil speed, rate of heat exchange, etc.) as well as variations in operation and maintenance practices are not taken into account;
5. Finally, the important effect of accessories and their component parts is not considered. For example, a single bushing may fail and lead the whole transformer to a catastrophic failure (end of life) and that may have nothing to do with the thermal aging illustrated in Fig. 1.

The work shown in the current paper is an attempt to overcome those pitfalls by integrating multiple existing engineering models (including thermal models and hot-spot estimation) into a larger data frame of statistical exploratory data analysis and Machine Learning algorithms for the assessment of the cumulative effect of multiple parameters on transformers remaining life. A R&D project was initiated in cooperation with utilities and transformer users in general to employ historical data of healthy and failed units and train multiple ML algorithms to try to best estimate remaining life.

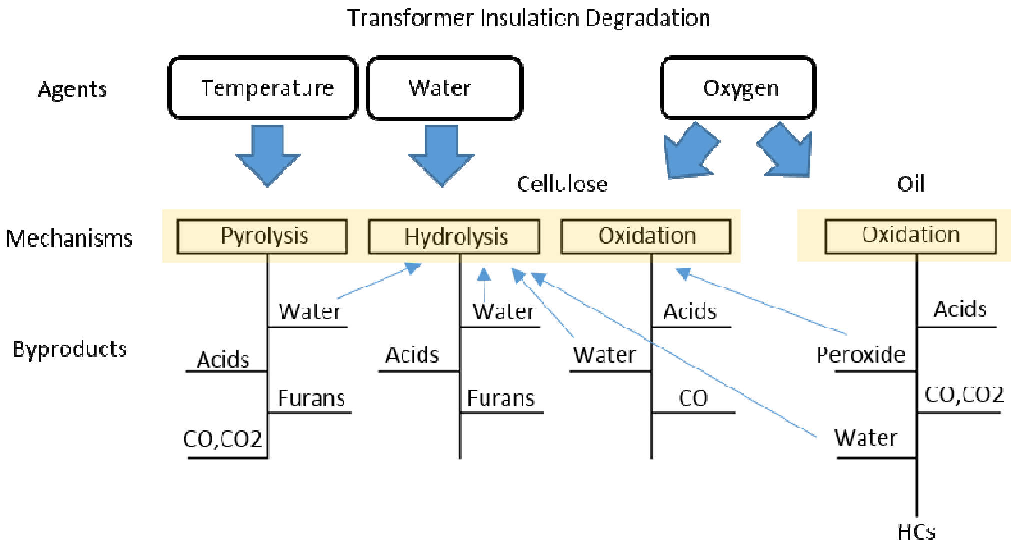

Figure 2 – Illustration of the complex degradation process of transformer insulation in the presence of the main aging factors (temperature, oxygen and moisture) as shown in [3].

2. Illustration of Initial Exploratory Data Analysis (EDA)

At this point only one utility (NES Power) and one industrial user have joined the data gathering effort. NES has about 170 power transformers in operation and a history of some 30 failed units since the 70's. Fig. 3 illustrates the distribution of legacy units, typically 69kV-161kV, 20-90MVA. Fig. 4 illustrates the units that failed since 1971 and the unit age at the time of failure.

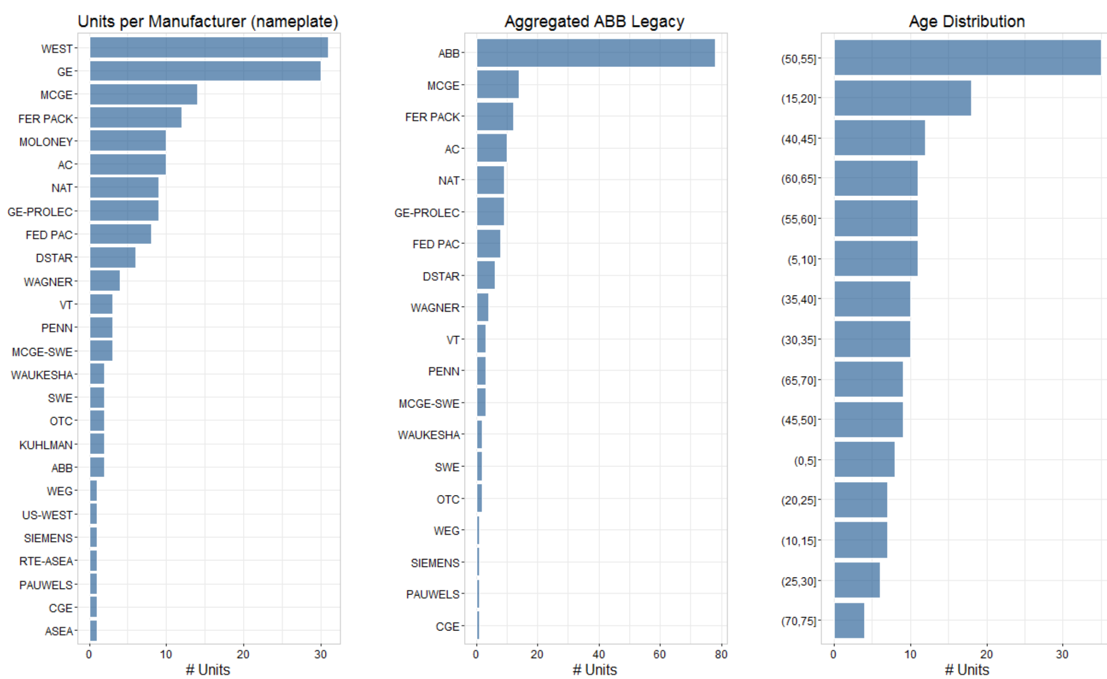

Figure 3 – NES legacy transformers and their manufacturers (left) with aggregated ABB incorporating several other suppliers that were acquired along the years (middle) and the distribution of the number of units per age group (right).

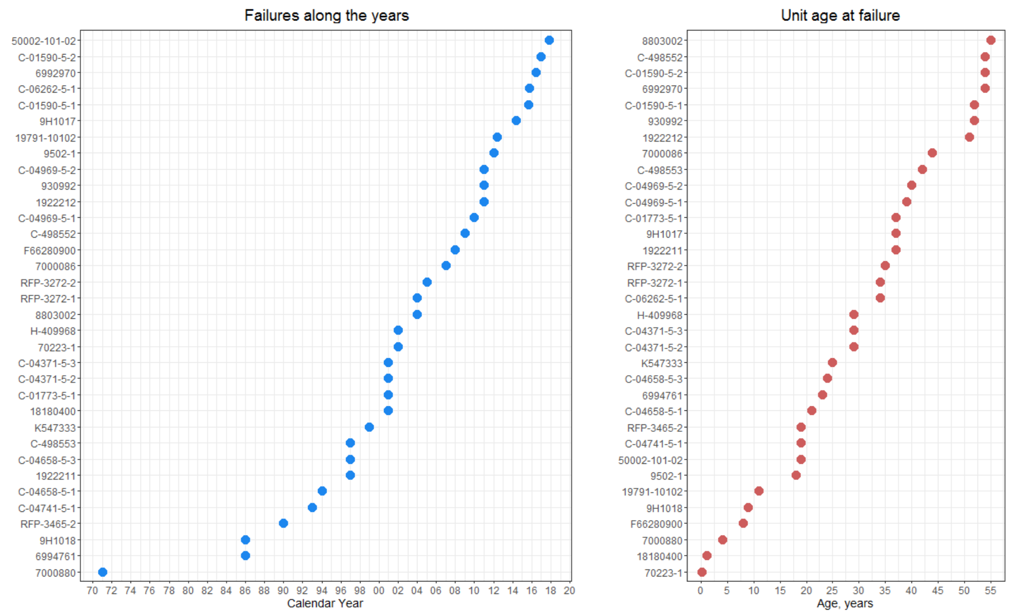

Figure 4 – Failures along the years from 1971 to 2018 (left) and transformer age at the time of failure (right)

2.1 Is Age a Factor?

Age is an important aspect that we would like to explore in our EDA since there are many investigators and even utilities to whom age is not a fundamental cause of failure or even a factor to consider. Although Fig. 4 confirms that failures do happen in all age groups, data shows, however, as illustrated in Fig. 5, that there could be significant statistical differences in behaviour of a given parameter for different age groups. The example shows the distribution of C2H2 (Acetylene) of all transformers per age group. Notice the singular behaviour of the estimated probability density functions for certain age groups. How much will that impact the probability of failure or the life of those units shall be part of the current investigation.

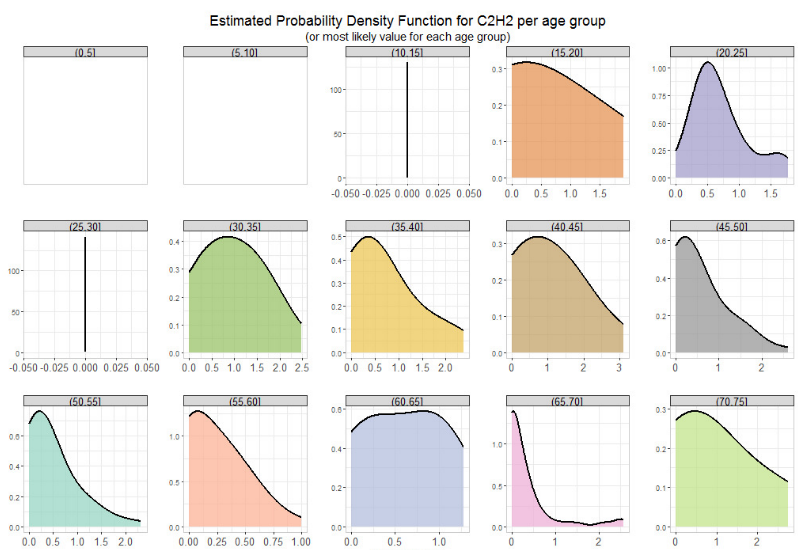

AgeGroup	Mean	Q90th	Max	Samples
(0,5]	0.00	0.00	0.00	52
(5,10]	0.00	0.00	0.00	90
(10,15]	0.01	0.00	1.00	135
(15,20]	0.33	0.00	80.00	323
(20,25]	4.01	7.00	59.00	173
(25,30]	0.02	0.00	1.00	140
(30,35]	3.18	2.00	303.00	260
(35,40]	3.09	1.00	236.00	241
(40,45]	12.27	1.00	1291.00	326
(45,50]	3.43	2.00	410.00	223
(50,55]	0.63	0.00	208.00	851
(55,60]	0.23	1.00	10.00	341
(60,65]	1.24	3.00	19.00	303
(65,70]	4.23	1.00	377.00	223
(70,75]	11.89	1.00	555.00	91

Figure 5 – Estimated probability density function for C2H2 (NES case) per age group. The table on the right of the picture gives the mean values, 90th Percentiles, max value and the number of DGA samples in each age group.

2.2 How bad does it need to be?

It is notorious that certain units may endure stresses in a much different way when compared to others that fail in the same condition. However, there are situations that even the seasoned transformer expert may find difficult to realize and fully understand why a unit has failed or even more interesting, why has it not. Figure 6 illustrates the case of all units in the study so far that have shown C2H2 levels greater than 1 ppm. The smaller chart on the right shows the maxima to all units represented on the left. The red dots show failed units as this would not surprise anyone, particularly as we move to the right hand side of the chart (very high levels of Acetylene). The white dots, however, show standing units or at least units that have not failed. For those levels of C2H2 this is remarkable. Further investigation is needed in order to identify the reasons for that as in, for example, Load Tap Changer (LTC) contamination of main tank oil, masking gas analysis. Notice that ppm scale is logarithmic. A very complex issue is the so called "pre-failure" parameters (all or most of the operational parameters and not only gases) and their impact on the actual failure mechanisms. This is also subject for a thorough investigation in this study.

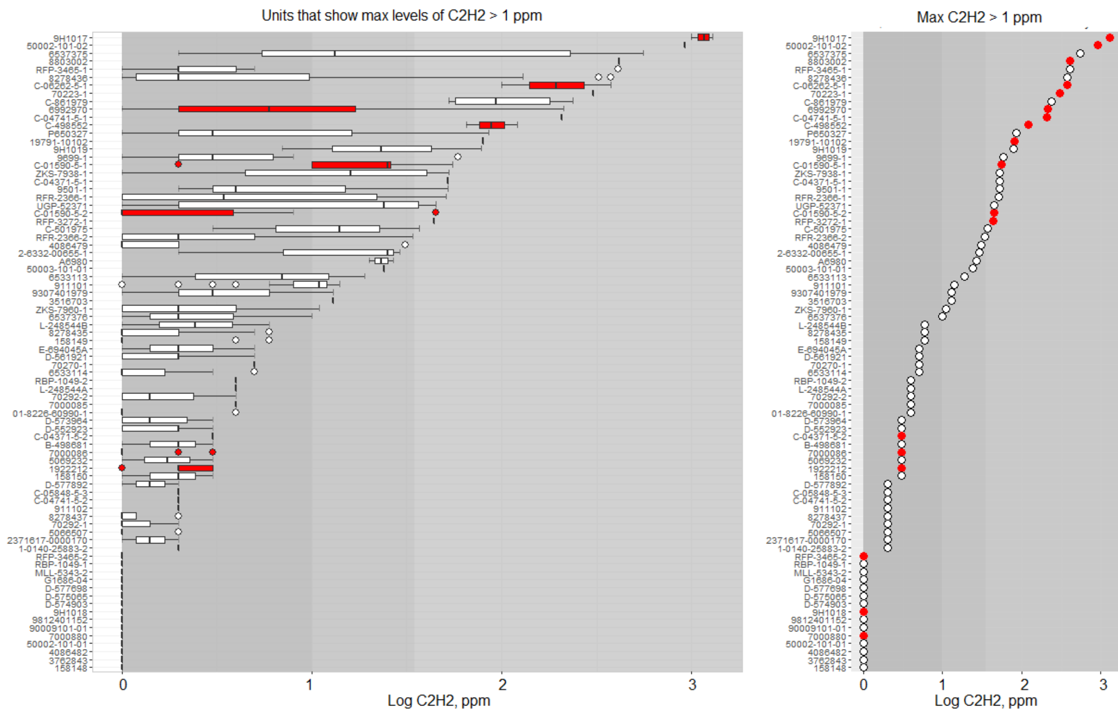

Figure 6 – (Left) Historical distribution of C2H2 for units that showed levels greater than 1 ppm at some point in their operation – red boxplots identify failed units; (right) - Maxima of C2H2 for units that showed levels greater than 1 ppm in (a) – red dots identify failed units. Shaded regions on the background show the limits for Condition 2 (above 1 ppm), Condition 3 (above 10 ppm) and Condition 4 (above 35 ppm) of the existing IEEE C57.104 Gas Guide [4].

Another interesting example illustrated with a different EDA technique (k-means clustering) is given in Figure 7. Here an algorithm was developed to analyse the relationship between carbon monoxide (CO) and carbon dioxide (CO2) that are, when at abnormal levels, frequently associated to paper involvement in an eventual fault. The algorithm is blind to any engineering interpretation and has no information about normal and abnormal levels of any gas dissolved in transformer oil. The cluster centroids were found purely based on mathematical similarity measures (Euclidean distance between points). The red and blue dots are the actual CO, CO2 pairs found in the NES dataset (historical levels of all transformers). An almost seamless boundary can be seen on the CO2 scales just above 2,500 ppm. Coincidence or not that is the Condition 1 limit for CO2 found in the IEEE C57.104 but the clustering algorithm does not know that. A second and perhaps more important observation is related to the failed units identified by the white triangles superimposed on the chart. A multitude of blue dots can be seen beyond the 5,000 ppm level for CO2 and both red and blue dots beyond 600 ppm for CO

(Condition 2 limit starts at 350 ppm and above in the IEEE C57.104) but the vast majority of the data coming from failed units (triangles) are concentrated in the lower region of both axes and absolutely no data from failed units is found on the extremes of either. This alone brings a lot of food for thought and will certainly be part of further and careful investigation.

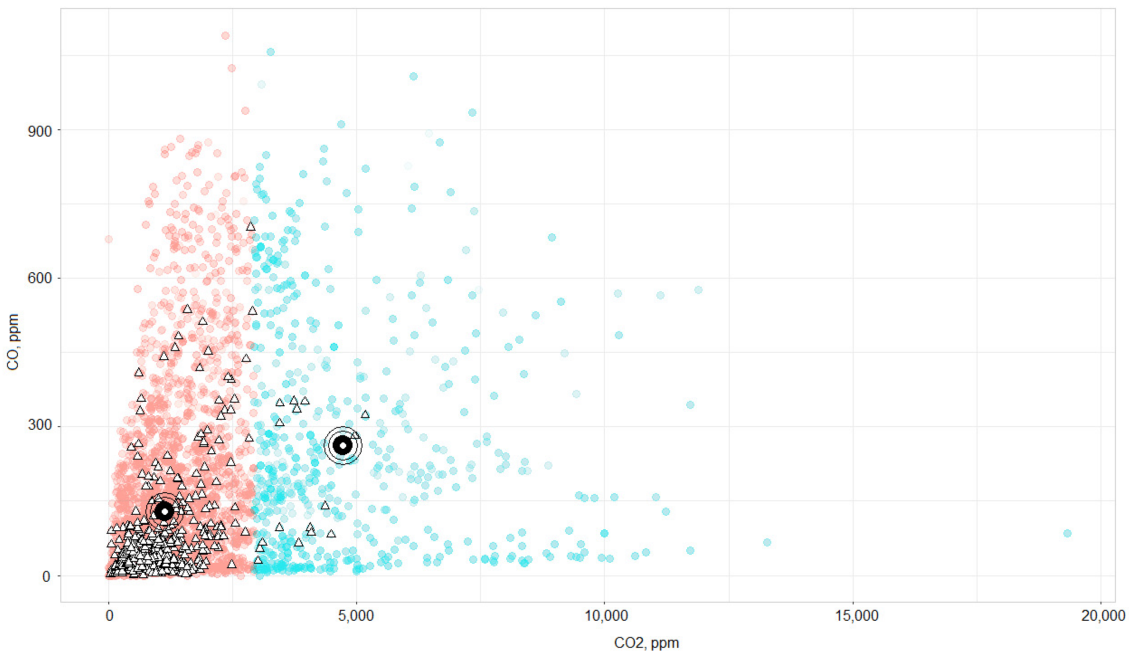

Figure 7 - Illustration of a k-means clustering (k=3) of CO2 vs. CO, showing the cluster centroids (large black circles), the data points belonging to each cluster (green, blue and red), superimposed by data coming from failed units only (white triangles).

3. Linear Regression

The first step into Machine Learning is almost always the use of simple linear regression to either look for trends or try to estimate one or more variables (dependent variables) based on the variation of others (independent variables or predictors).

At this stage a simple trend analysis algorithm was created to establish first if a given parameter (say a given combustible gas) would show positive correlation with time. Failed units were identified so that their trends could be visual. If a positive correlation was found then separate those cases which also had good coefficient of determination ($r^2 \geq 0.5$) and plot them in descending order of slope (higher slope = higher gas generation rate and hypothetically higher the risk of failure). The procedure was then repeated to multiple parameters. Figure 8 illustrates the case of the evolution of C2H6 (Ethane) with time for all units in the so called small multiples plot for the NES dataset. Notice the lines in red representing failed units. Each small panel brings the unit id at the top.

Figure 9 illustrates only the units that showed positive and significant correlation with time and that showed trend at significant gas levels (positive and good correlation at very low gas levels were neglected). The red tods in some of the facets indicate failed units. Here again we see the need for a multivariate approach to the identification of possible failure candidates since the units with higher gas generation rate are not necessarily the ones that failed and there are examples in the figure that illustrates the opposite: failed units with positive but relatively low gas generation rate, as for example, the very first panel from bottom to top on the right hand side. Some failed units even showed negative trend for C2H6 as illustrated in Figure 8, fifth panel on the right hand side, top to bottom.

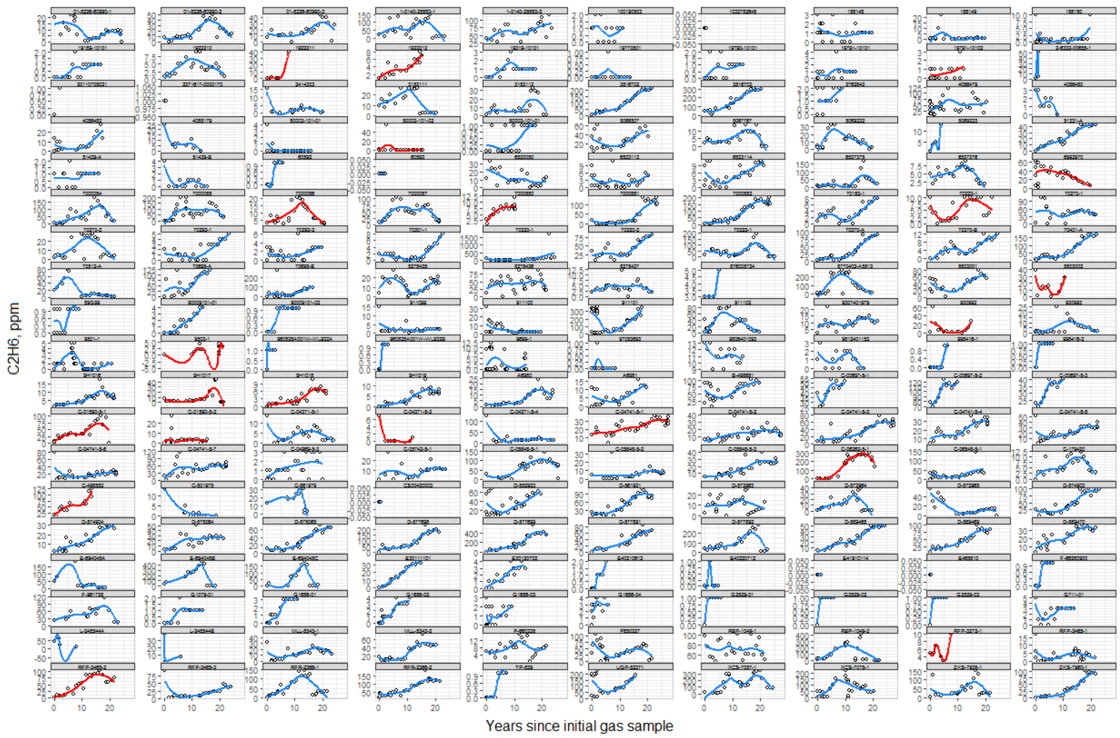

Figure 8 – Small multiples showing evolution of C2H6 with time to all 190 units in the dataset. Red lines show failed units. Notice that some units in red show descending C2H6 rates (negative slopes).

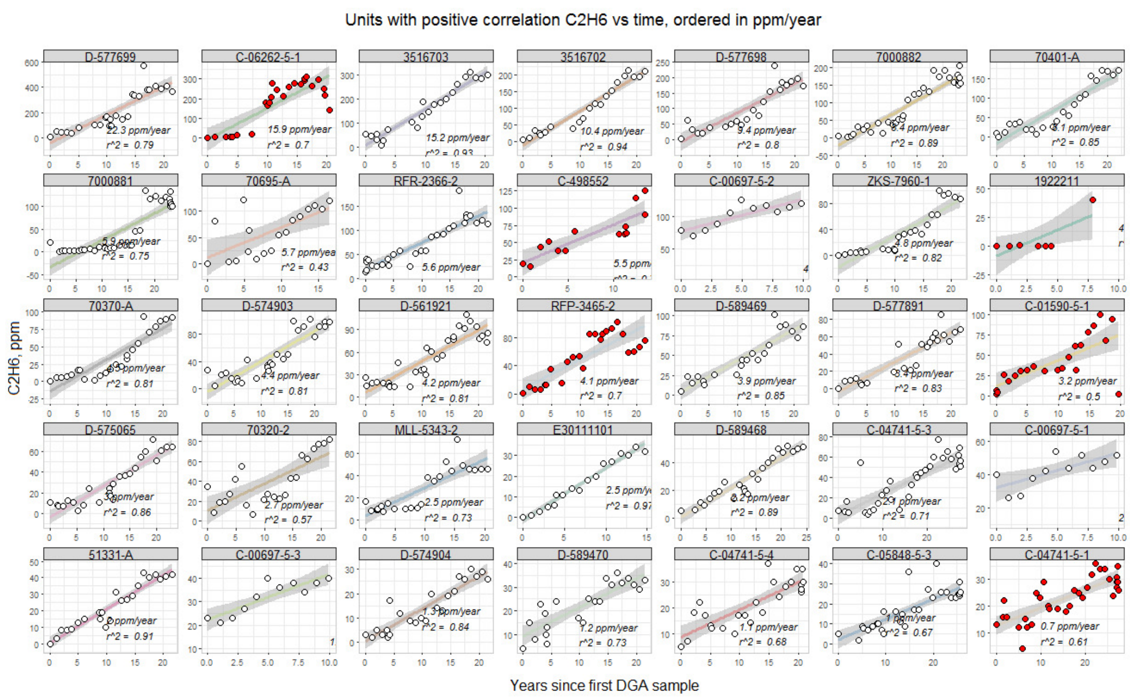

Figure 9 - Units that showed significant combustible gas trend (C2H6) with good positive correlation with time (coefficient of determination $r^2 \geq 0.5$) ordered from the highest gas generation rate (top left = 22.3 ppm/year) to the lowest gas generation rate (bottom right = 0.7 ppm/year); all units showing significant amounts of C2H6. Red dots identify failed units. Notice that failed units are not necessarily those with the highest rate of gas formation.

4. Discussion and Future Work

The examples above are simple illustrations of the many aspects involved in developing a comprehensive engineering based statistical approach to try to estimate power transformers life expectancy based on a large number of parameters that certainly have a complex relationship among them. Aging and the consequent weakening of the solid insulation is a very complex phenomenon, as illustrated in Figure 2. In practical terms there is a reduction in the paper tensile strength and its Degree of Polymerization (DP) that at certain low levels (say < 200) may lead to mechanical rupture of the solid insulation and eventual breakdown. Notice the negative effect produced by some of the aging by-products in a closed loop feedback system. Although the phenomenon is reasonably well understood and the effects well studied in the laboratory, the actual quantitative analysis of the impact of each individual contributing factor and their aggregate effect on aging is extremely complex and difficult to determine in real life transformers for the reasons already mentioned above.

The paper shows the initial stage of our statistical analysis, in preparation for the use of Machine Learning algorithms, of a large number of historical parameters from a small fleet of healthy units and failed transformers. The continuation of the current investigation with the application of more sophisticated Machine Learning algorithms in the future for better life estimates and predictions is heavily dependent on the availability of more data and data of good quality, particularly involving failed units. The R&D team hopes to find more support in the industry and that other transformer users may join the effort in the near future.

BIBLIOGRAPHY

[1] Cigre SCA2 Paris Session Paper (Best Paper Award), Machine Learning in Support of Transformer Diagnostics, Sept. 2018

[2] IEC60076-7/2018-01, Loading guide for mineral-oil-immersed power transformers

[3] Working Group SCA2.24, Thermal Performance of Transformers, Cigre Brochure 393, 2009

[4] IEEE C57.104/2008, Guide for the Interpretation of Gases Generated in Mineral Oil-Immersed Transformers

[5] IEEE C57.106/2015, Guide for Acceptance and Maintenance of Insulating Mineral Oil in Electrical Equipment