Anode Spike Detection Using Advanced Analytics and Data Analysis

Arthur Martel
Process Control Engineer
Rio Tinto, St Jean de Maurienne, France
arthur.martel@riotinto.com

Abstract

Anode spikes crises have a deleterious effect on current efficiency to the point of jeopardizing smelter operation. All the time spent with a spike under an anode is a period where current is lost, and the longer the spike remains present, the more serious are the mid-term consequences for the reduction process. It is therefore most important to detect these spiky anodes as early as possible and remove them from the pots. A new tool based on anode current measurements, combined with machine learning, has been developed and tested. It is an effective way of detecting many of these spikes, usually a few days before they become obvious. This article describes the development of the tool and the first results obtained on industrial cells.

Keywords: Aluminum electrolysis, anode spikes, machine learning.

1. Introduction

The effects of anode spikes in aluminum electrolysis are well-known. Examples are the disruption of operations, the thermal and chemical imbalance of pots and the loss of current efficiency; and there are many more. Especially after a spike crisis, some of these consequences can degrade operations for weeks or even months after the faulty anodes have been changed. Studying the reduction in alumina feeding in pots before spikes are spotted clearly shows the loss of current efficiency involved by such incidents, as seen in Figure 1.

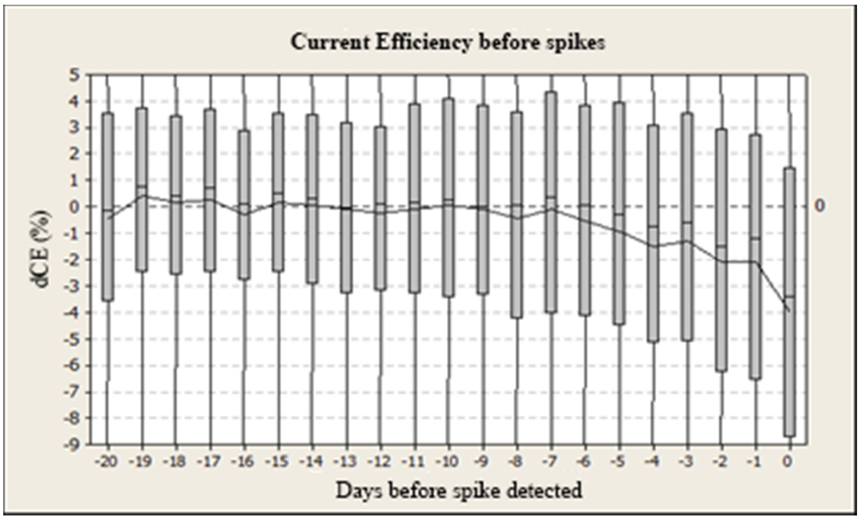

Figure 1. Current efficiency over the period before a spike is detected.

This indicates how important it is to react as fast as possible, since the loss of production (and the extent of the consequences, especially on thermal and chemical balance) is proportional to the time spent with a spike on the pot. Effects are even bigger when the spike is fully developed, just before the faulty anodes are changed, with more than 3 % reduction in current efficiency. A study has been carried out to see how data analysis, through machine learning models, can help process teams to spot such spikes as early as possible.

General Approach & Measurements

This seemed to be a good opportunity to make use of anode current measurements [1]. Not only can they be expected to contain information enabling spike detection, they are also available in sufficient detail to make it possible to locate the faulty anode(s), instead of just flagging the pot itself.

In some of Rio Tinto's R&D pots anode current measurements are made at a rate of 1 Hz, much more frequently than with the usual electrolysis pot instrumentation, so these were the data that were analyzed in order to provide the required tool.

2. Statistical Modeling

2.1. Hidden Markov Models

The approach that was chosen for this study was to build a classifier intended to distinguish "healthy" anodes – those without spikes – from anodes having a spike. Anode currents are measured continuously, but for the analysis they are considered as discrete time series. It is difficult to predict when a spike will be detectable, even if anodes seem more likely to develop spikes at certain periods in their operating life, as can be seen in Figure 2. Most of the spikes are spotted either around the first third of the anode's lifetime or else right at the end, when it is in any case due for changing. This pattern may be technology-dependent, since the spikes in the pots that were studied were mainly on the edges of the anodes, probably starting to develop right after anode changing.

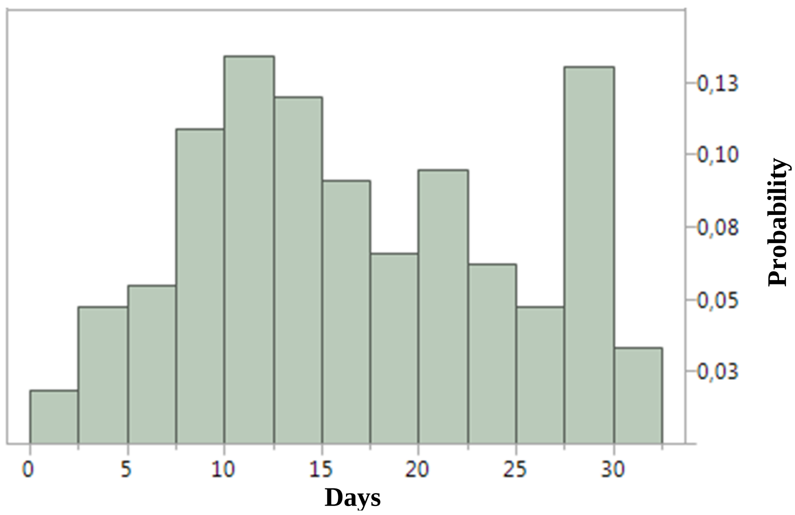

Figure 2. Distribution of the number of days before a spike is spotted.

These properties of the data being analyzed led us to consider using Hidden Markov Models (HMM's) [2]. HMM's are statistical models that allow analyzing sequences of observations, and are designed to identify state changes - transitions in these observations - enabling pattern recognition. They were firstly used in speech recognition, where words are successions of sounds of different properties and variable durations.

A similar approach can be applied to the analysis of the anode current, where the amperage signal over an anode life (or features taken from this signal) is the sequence, and the HMM converges to detect patterns in this signal and the transition from a normal anode to a spiky one - although the model can actually detect finer state transitions than this.

2.2. Data Pre-Processing

As in any data analysis, information has to be preprocessed according to the needs of the study and what is required in the dataset. This is sometimes ignored in algorithms such as Deep Learning, where features are in fact generated by the model itself, but this requires a much greater mass of data than that available in a spikes database. Consequently, two main features have been extracted from the anode current signals, according to the knowledge we have of the phenomenon; these are described in Section 3.3 below.

The process of spike development comprises of three steps, as follows:

- Step 1 (Figure 3). An area of the anode is less conductive, or even insulated. This could be caused by dirt and bath pollution, the quality of the anode itself or to thermal imbalances in the anode, especially just after changing (leading to spikes or deformation on the sides of anodes).

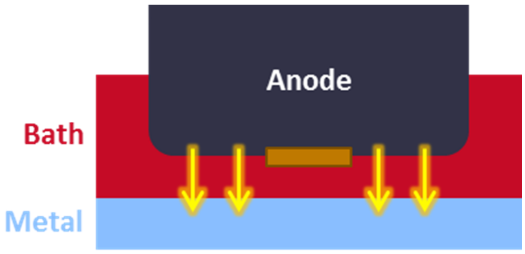

Figure 3. Less- or non-conductive area under an anode as the origin of a spike.

- Step 2 (Figure 4). The spike begins to appear under the anode. At this stage waves in the metal pad could well reach the growing spike, thus creating a temporary local short-circuit, or at least increasing local variations in anode current density.

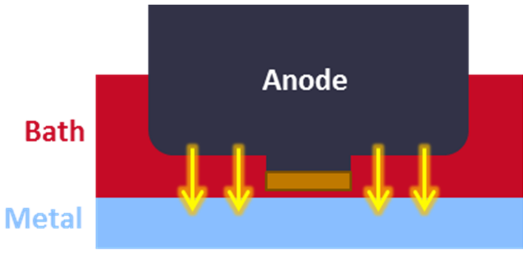

Figure 4. Beginning of a spike development.

- Step 3 (Figure 5). At this stage the spike is fully developed and its effects are at their most severe. The spike extends right into the metal pad, causing a permanent short-circuit. The short-circuit current represents a pure loss of current efficiency, since it is not used for electrolysis. A local hotspot is generated, the current distribution is disturbed and the anode beam can even rise, causing thermal and chemical imbalances.

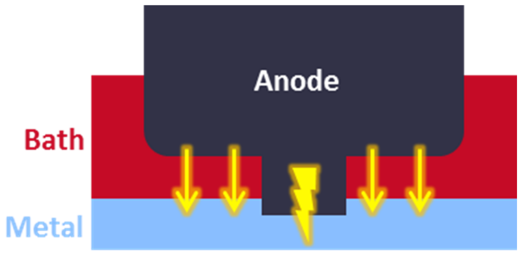

Figure 5. A fully developed spike.

The raw data have to be pre-processed in a way that makes it possible to identify these steps, by calculating features expected to vary throughout this process. For example, the range of the anode current can be considered, as well as its short-term level of noise. Nevertheless, just using the level of current is not statistically differentiating enough, since all anodes don't actually stay at a constant amperage level. Even if the global average current for all anodes shows a typical pattern (depending on anode slots and operations cycle), as shown in Figure 6, this variation in a normal anode is actually strongly dependent on its position in the pot, and on the shape of the metal pad.

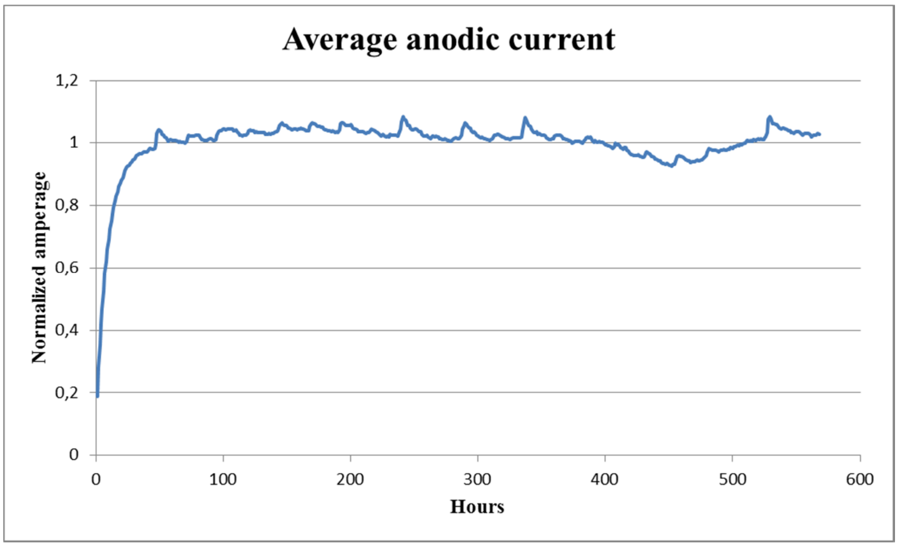

Figure 6. Average anode current.

The same figure drawn for each anode position actually includes much more variation, and precise patterns can be more diverse. Figure 7 shows the average current calculated for two anode positions over a dozen anode cycles. The peaks that can be observed are linked to operations cycles, and are actually recurrent.

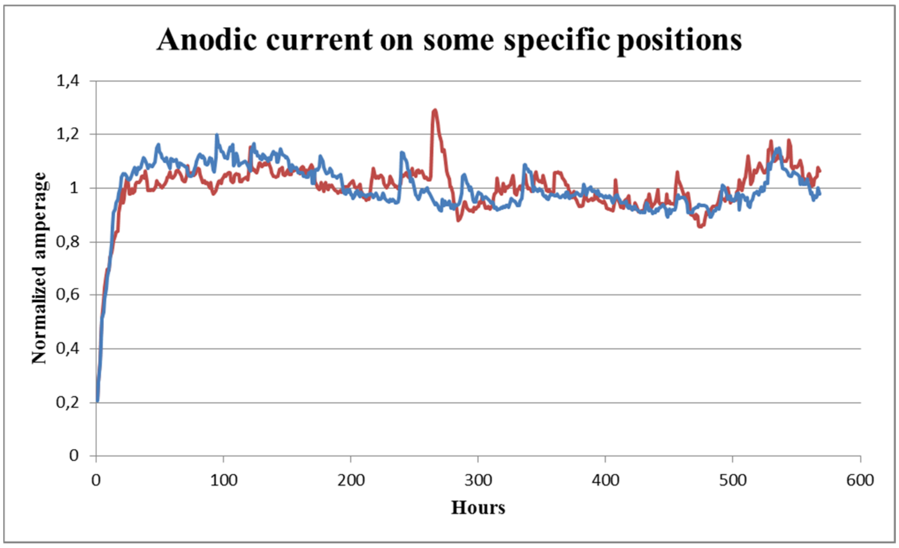

Figure 7. Anode current for two anode positions.

In order to detect spikes, small variations in anode current have to be spotted, which may be lower than the overall range of anode current (even without considering anode current pick-up), and measured amperage is thus much more statistically significant when compared to a reference calculated for each anode position as shown above. Standardization of this kind can be applied to other features too.

2.3. Machine Learning

In order to train the classifier, sequences of the time-series of anode current observations are built with the different features. The results shown below are from models that use only 2 features:

- Anode current compared to its reference (calculated for previous anodes without spikes, at the same position)
- Standard deviation of high-pass filtered anode current.

These sequences start when anodes are installed in pots, and stop when they are removed (or when a spike was detected if earlier). The classifier works as a competitive classifier, so two models are actually trained. The first model is trained on the group of sequences of healthy anodes, which had not reported as having a spike, while the second one is trained on sequences of anodes where a spike was spotted. This makes it possible to calculate for each sequence the likelihood of presence of a spike, and the likelihood of absence of a spike. Comparing these two values (with a possible threshold) is a way to classify the sequence. Figure 8 shows the comparison of these two scores. Healthy sequences are black dots, and anodes with spikes are red crosses. The green line is the limit of equal likelihoods. Above the green line, a sequence is more likely to correspond to an anode with a spike, and under the green line it is more likely to be a healthy anode. Please note the negative axis scales: these are logarithmic probabilities.

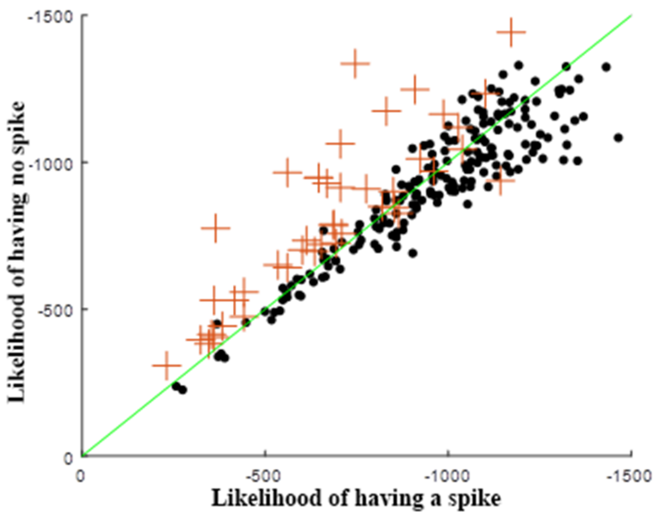

Figure 8. Likelihood scores for healthy and spiked sequences.

Cross-validating models with this classification strategy makes it possible to try to get the best parameters and variables. Yet another step is required to assess model performance: calculating these likelihoods at every time step of the sequences gives a more realistic idea of the actual occurrences of alarms. Some sequences can indeed lean toward one side before coming back to the other side of the decision limit at the end, thus increasing the number of alarms compared to what can be measured on full sequences. The results of this calculation can be seen in Figure 9, the red lines indicating anodes with spikes. This gives better information on how the green line can be translated to adjust the false versus true alarm rates.

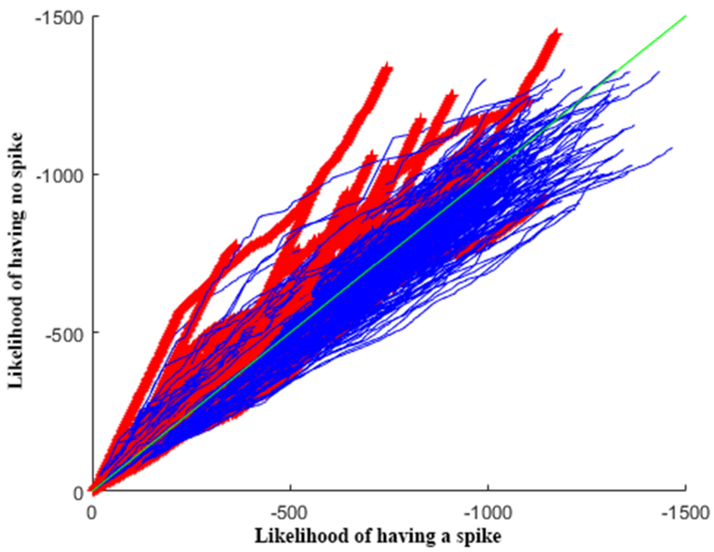

Figure 9. Online estimation of likelihoods.

2.4. Operational Choices

At this stage, the models are trained and test sequences are scored over their entire duration. Armed with this information, it is possible to calibrate the decision threshold that raises alarms, in order to get a compromise suiting local needs. In such cases, since the effects on operations will be felt on raised alarms, indicators about the positives (detections) seem more relevant than some other indicators that may be more common in machine learning algorithms assessment (accuracy, specificity, area under receiver operating characteristic curve, etc.). In this case, the following indicators have been chosen to assess the performance of the models:

- Precision, i.e., the proportion of alarms that detect an actual spike (Equation (1))

$$Precision = \frac{\text{Number of detected spikes}}{\text{Total number of alarms}} \tag{1}$$

- Recall, i.e., the proportion of spikes the algorithm succeeds in detecting (Equation (2))

$$Recall = \frac{\text{Number of detected spikes}}{\text{Total number of spikes}} \tag{2}$$

- Time-lapse between prediction of a spike and its becoming obvious.

The drawback of *Precision* in particular is that it is dependent on the actual proportion of spiky anodes. If spikes remain exceptional, even the tiniest rate of false positive will dramatically decrease this indicator. Consequently, the contribution of this indicator is more relevant during spike crises than in periods of normal operation. Figure 10 shows the balance, in one modeling result, between *Precision* and *Recall*. Raising the decision threshold gives better *Precision* (alarms are raised only on sequences that clearly look like spikes) but tends to reduce *Recall* (the most ambiguous spikes are excluded from the alarms), while reducing this threshold will invert these effects. If sample sizes are small there may be discontinuities, as in Figure 10.

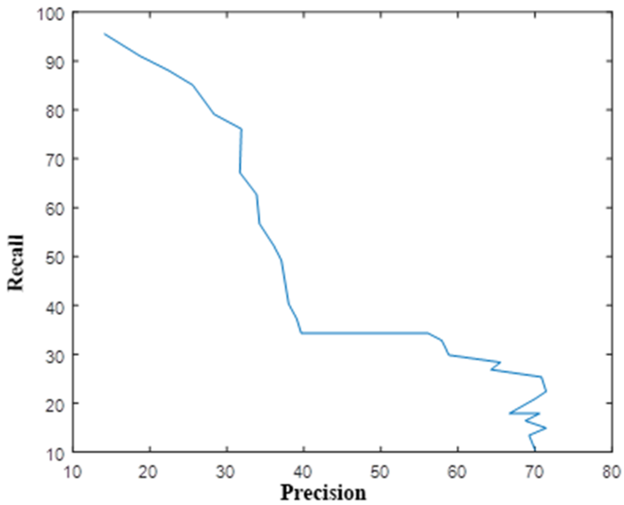

Figure 10. Balance between precision and recall, depending on decision threshold.

In this case, depending on staff availability, the estimated cost of spikes and that of checking suspect anodes that turn out to be healthy, we could set a target - for instance, near 60 % precision and 35 % recall - or else another one near 35 % precision and 65 % recall.

The approach described so far gives the main steps required to get such a system working. Nevertheless, such algorithms can be improved in many ways:

- Structure of models (parameters of HMM),
- Filtering the data to eliminate outliers,
- Refinement of features (filters, thresholds, standardizations, etc.),
- Addition of new features (bath temperature, chemistry, orders to raise/lower the beams, pot state clustering, etc.),
- Determining models' learning by possible operations, incidents or other events.

Every improvement leads to new performance data with which to assess its value.

3. Data driven management

Beside purely mathematical and algorithmic considerations, another aspect that might be underrated and forgotten is management and organization around data usage. A whole chain must be set up and maintained, involving:

- Measurements (sometimes manual),
- Data acquisition and data gathering strategy,
- Algorithm development,
- Including algorithms in process control industrial tools, or providing access to their results in an easily usable way,
- Training users,
- Considering eventual changes in managing operations and in operating practices
- Feeding back the algorithms.

All these steps require the understanding of how a data-oriented project cannot just come down to manipulating data and building algorithms. These may be the core of it, but it cannot exist without some prerequisites, nor succeed without aligning the construction of models with operations management. This can be required very early in the process, in this case for example in properly flagging spikes, since flagging is mandatory in most of the machine learning algorithms.

4. Conclusion

This proof of concept shows how machine learning can help analyze data and statistically detect incidents provided the relevant data are available and given the flagging of incidents in operational logs. In applications such as this, given the low number of samples, a good knowledge of process and data is required to target measures and build relevant features. This process can be time consuming, especially when using data that were not originally intended to be used in this way. Nevertheless, gathering data as soon as possible, building up bigger and more accurate databases and getting used to such tools seems to hold out promise for use in similar cases.

5. References

1. Jeffrey Keniry and Eugene Shaidulin, Anode signal analysis – The next generation in reduction cell control, *Light Metals* 2008, 287-292.
2. Walter Zucchini and Iain L. MacDonald, *Hidden Markov Models for time series: an introduction using R*, Chapman and Hall/CRC (2009), ISBN 9781420010893.