AI-Driven Anode Effect Prediction in Hindalco Renukoot Smelter

Rupesh Pandey[1] and Faraz Khan[2]
1. Assistant Vice President
2. Data Scientist
Hindalco Industries, Renukoot, India
Corresponding author: rupesh.pandey@adityabirla.com
https://doi.org/10.71659/icsoba2025-al052

Abstract

The Hall-Héroult Process is used to produce aluminium from alumina using the electrolysis process in an aluminium reduction pot. Alumina is usually fed into the pot at regular intervals, using a point-feed mechanism. At very low alumina concentration (< 2 %) cryolite decomposition occurs and perfluorocarbon (PFC) gases evolve in the process. These gases form a highly electrical resistive layer underneath the anode, which suddenly increases the pot voltage, resulting in an anode effect (AE). High anode effect frequency (AEF) and anode effect duration (AED) may lead to reduced anode life, high heat generation, crust collapse, ledge melting, low current efficiency, high fluoride emission, and the possibility of pot leakage. Therefore, high AEF and AED contribute significantly to PFC emission and global warming. Apart from the environmental challenges, AEs significantly increase specific energy consumption and affect pot life. Aluminium smelter industries are aggressively examining ways to reduce the number of AEs as well as their duration.

To mitigate these challenges, an in-house AI-driven AE prediction model has been developed, utilizing real-time pot data from the control system to forecast AE occurrences up to one hour in advance. This model, seamlessly integrates with an in-house software system, enables proactive control measures, effectively reducing both anode effect frequency (AEF) and anode effect duration (AED). With an accuracy ranging from 85 % to 90 %, the system has significantly minimized manual interventions, reducing workplace hazards, improving operational stability, and enhancing overall energy efficiency. Additionally, by optimizing process control, the model aligns the smelter's performance with global benchmarks, ensuring sustainable aluminium production.

Keywords: AI-driven model, Aluminium smelting, Anode effect prediction, Energy efficiency, PFC emissions reduction.

1. Introduction

The Hall-Héroult Process is used for the extraction of aluminium (Al) from its oxide, alumina (Al_2O_3). In this process, alumina is fed into the system through point feeders, dissolving in molten cryolite (Na_3AlF_6) for electrolysis to liquid aluminium (Al) and carbon dioxide (CO_2). The Hall-Héroult process operates in the temperature range of 940–980 °C and typically produces aluminium with a purity between 99.5 and 99.8 %. Additives such as the fluoride salts AlF_3, MgF_2, CaF_2, and LiF are mixed with cryolite in order to optimize the process. These additives reduce the liquidus temperature and increase the efficiency of the electrolysis of alumina [1].

The Hall-Héroult Process is highly energy-demanding, wherein electricity constitutes 30–40 % of the total cost of production. A schematic showing the process is given in Figure 1. Modern aluminium smelters achieve current efficiencies (CE) nearing 94 %, while best-practice benchmarks go up to 96 % [2]. Maintaining high CE requires strict control of the alumina concentration, keeping it within the optimal range, along with rigorous compliance to standard operating procedures (SOPs).

In earlier days, when automatic feed control systems did not exist, the operator would detect the anode effect (AE) when a spike in the cell voltage occurred. This signal alerted the operator to manually add a certain amount of alumina into the cell. The interval until the next AE occurred allowed calculating the adequate feeding rate. For example, if AE is too early, increase the alumina dosage; if too late, decrease it. Manual termination of AEs usually occurred as either scraping by iron rakes or by green wooden poles. As a result of the above, AEs cause melting of the side ledge and localized sludging. More critically, AEs lead to the emission of perfluorocarbon (PFC) greenhouse gases, mainly CF_4 and C_2F_6 which possess very high global warming potentials [3–7].

With the advent of the computerized automatic feeding systems, computer controlled dosing facilities, and automated AE termination protocols, the aluminium industry has considerably reduced both occurrences and durations of anode effects. Currently, most aluminium smelters have adopted point-feed technology for the alumina introduction into the electrolytic cell. This system is generally provided with a cylinder-type volumetric feeder, which works side by side with a crust breaker. Normally, alumina is added in discrete batches, whereby each feeder delivers its precisely measured quantity of 1 to 2 kilograms into the bath through a designated feeder hole at predetermined time intervals. The initial feed control strategies were based on cell voltage which reflects several factors such as alumina concentration, anode-cathode distance, electrolyte conductivity, and the operating anode current density.

Renukoot Smelter has been in existence since 1962, and over these years, it has continued to modernize pot designs including upgrades of various registers and control systems. Such advancements have ensured the smelter remains competitive with state-of-the-art technologies. Renukoot is producing at 410 kt/a and achieves current and energy efficiencies comparable to modern smelters. The process control systems at the smelter have improved significantly from manual alumina feeding to an automated, resistance-based feeding system, referred to as dynamic feed, which includes hardware and design improvements. The control system is further enhanced by incorporating the AI-based models. This paper is about development of AI-based AE prediction model which enables early detection of AE.

Hindalco Renukoot

- Potline amperage = 70 kA
- Metal production: 530 kg/pot·day
- Number of pots: 2138 in 11 potlines
- Current density: 0.94 A/cm^2
- Metal production: 410 kt/a

Figure 1. Brief overview of aluminium smelter at Hindalco Renukoot.

1.1 Literature Review

In terms of practical improvements and more research into alumina feeding as well as anodic effects, there have been no recent advancements. Haupin and Seger [3] proposed a variety of methods for predicting AE initiation, including hysteresis analysis in the V-A curve, monitoring voltage increase rate, detecting high-frequency electrical noise, measuring acoustic emissions, and using pilot anodes. They also enumerated remedies for minimizing AE frequency such as keeping pot lining and electrical connections, maintaining alumina quality, maintaining feeder equipment, achieving uniformity in anode current density, and maintaining a suitable ledge profile. Thorhallsson [8] suggested that improved maintenance practices and increased attention to pot tending would reduce the incidence of anode effects. He proposes enhanced maintenance practices, increased attention to pot tending to lower AEF, and a gradual downward adjustment of the anode beam with timed intervals between each step. Additionally, an optimized feeding strategy during termination is recommended to shorten AED.

Fardeau et al. [9] reported that a good control of bath height is an essential element to reduce AEs. Mulder et al. [7] identified the root causes of the AEs classified into three areas: pot controller settings, feed strategy, and combination of maintenance practices and alumina quality. Individual anode current measurement (ACM) has brought forth models and techniques that may be used in predicting AE [10–13]. In the context of prediction models based on classic process data, da Costa et al. [14] and Wilson et al. [15] developed univariate threshold-based approaches. This resulted in a lesser lead time (in seconds) which was defined as the time difference between the model's detection and the process control system's detection. Y. Zhang [16] applied wavelet packet transforms to analyse pot resistance for AE prediction. Majid [17] employed multivariate statistical techniques for AE prediction. Z. Zhang et al. [18] developed a machine learning model achieving a significant lead time of 30 minutes and an F1 score of 99.8 %. Chen et al. [11] introduced a collaborative two-dimensional forecasting model which provided a lead time between 20 and 40 minutes with an accuracy of 95 %.

In general, looking at all the proposed methods, none of them seem to have been validated on sufficiently large real-world datasets or tested in a live operational environment. What is more, a number of studies do not use comprehensive evaluation metrics, such as confusion matrices or F1 scores, making it quite problematic to assess their effectiveness. Hence, a predictive model was developed based on supervised machine learning. It aims to forecast anode effects with a 20-minute forecasting horizon for timely operator intervention (e.g., manual alumina feed or process intervention). The major functional feature includes prediction every 20 minutes per pot: it guarantees real-time AE risk status for the plant control room, enabling early action within the 20-minute window before a possible AE event [19].

1.2 Anode Effect

The anode effect is a situation which starts when the electrolyte no longer wets one or more of the anodes and subsequently, the voltage rises beyond its normal value. The primary reason for the anode effect occurrence is higher concentration overvoltage [4]. The decrease in the concentration of alumina results in the increased surface tension of the electrolyte thus forming larger gas bubbles on the anode. In this way, the increase in current density on the active parts of the anode enhances the overvoltage to the point where discharge of fluoride ions, the next most easily oxidized ions, begins. This condition results in formation of perfluorocarbons (PFC) gases. Therewith, a continuous gas film is formed between the anode and electrolyte, resulting in a very high cell resistance.

The phenomena lead to a continuous formation of PFCs gases such as CF_4 and C_2F_6, which escape from the cell as shown in Figure 2. The produced gas at the anode reduces from practically all CO_2 before AE to 10–20 % CF_4 by volume, 1–3 % C_2F_6, 2–10 % CO_2 and the balance CO, during AE [3].

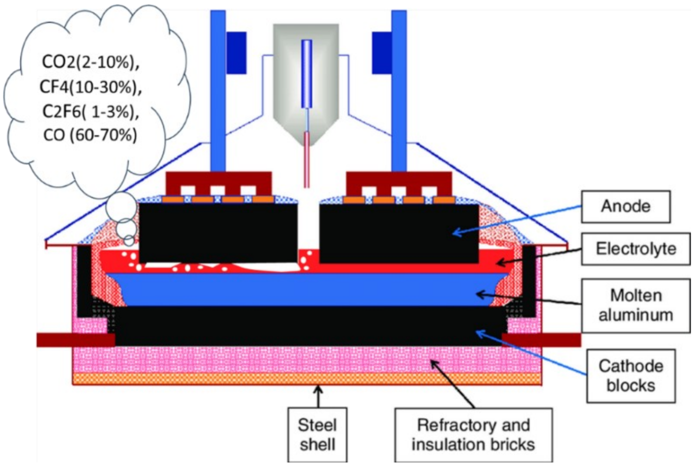

Figure 2. Formation of gases during the anode effect.

2. Advancements in Pot Controller Technology at Renukoot

The advent of pot controllers has set in motion a groundbreaking change in our aluminium electrolysis operations. Advanced pot controllers backed with new technology have enabled us to track key process parameters with unprecedented diligence. Advanced Process Controllers (APC) have been developed by in-house talent, the evolutionary journey of which is shown in Figure 3.

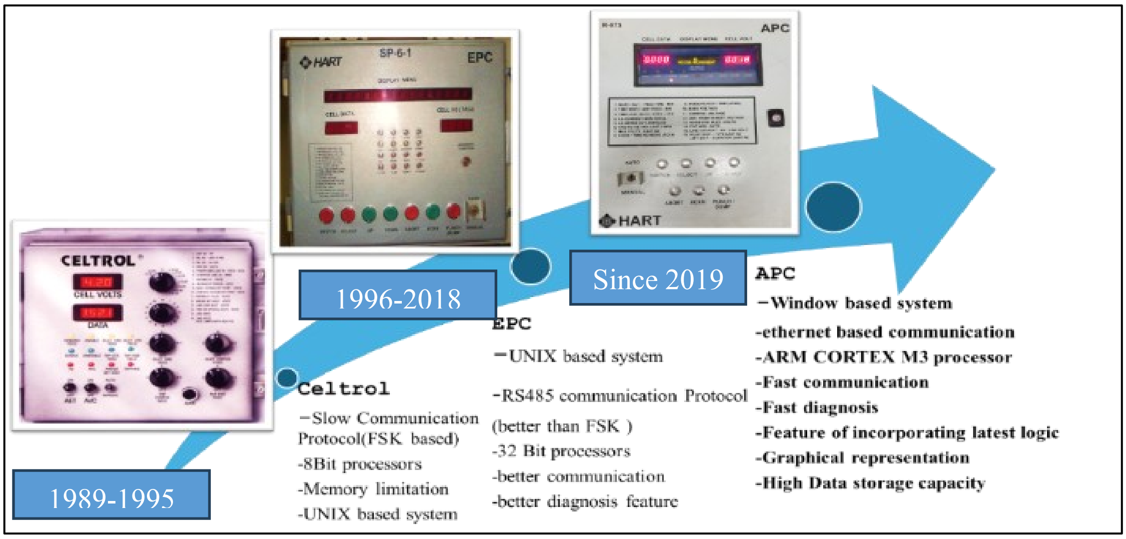

Figure 3. Evolution of pot controller at Renukoot.

The APC is newly designed, which is an upgraded version of the old Electronic Process Controller (EPC). It collects data at high frequency, is capable of complex calculations, pot trends display, and high-frequency data availability for analysis. With high frequency and precision in monitoring, adjustments can be made quickly to avert anode effects and maximize the general process efficiency. Predictive maintenance strategies are therefore complemented by continuous improvement. Implementing Windows based pot controllers and faster data acquisition has enabled for development of Machine Learning-based AE prediction model.

3. Data Selection and Preprocessing

The dataset used in this study was originally recorded as raw sensor data from aluminium electrolysis pots, collected continuously with high temporal resolution. To ensure data quality and relevance, a series of preprocessing steps were applied:

- First, only a subset of columns deemed relevant to the analysis was selected. These columns included measurements such as noise levels, voltage readings, control signals, and temporal markers. This selective approach reduces computational load and focuses on variables hypothesized to influence acoustic emission events.
- Because several measurement columns stored values in a hexadecimal-like format, a custom transformation was applied to convert these values into meaningful numerical scales. The transformation involved reversing the byte order and normalizing the resulting integer values by a fixed factor (800), effectively converting raw sensor readings into interpretable units.
- A binary acoustic emission indicator was computed based on threshold criteria combining noise, search h without noise, and voltage levels. Specifically, an AE event was flagged if either the voltage exceeded 10 V or a combination of zero noise and zero search coincided with voltage above 10 V.
- To focus on valid operational data, records with voltage below 3.75 V were filtered out, removing low-activity periods likely irrelevant for AE analysis.

For each pot, AE events were used as temporal anchors to extract a 30-minute window before and after each AE occurrence, ensuring the capture of pre-and post-event behavior. When AE events occurred in close succession, the window was truncated at the next AE event to avoid overlap and redundancy. This is a unique procedure to reduce data imbalance (between no AE and AE classes) as shown in Figure 4. Data is balanced from (99.8–0.2) % to (74.6–25.4) %.

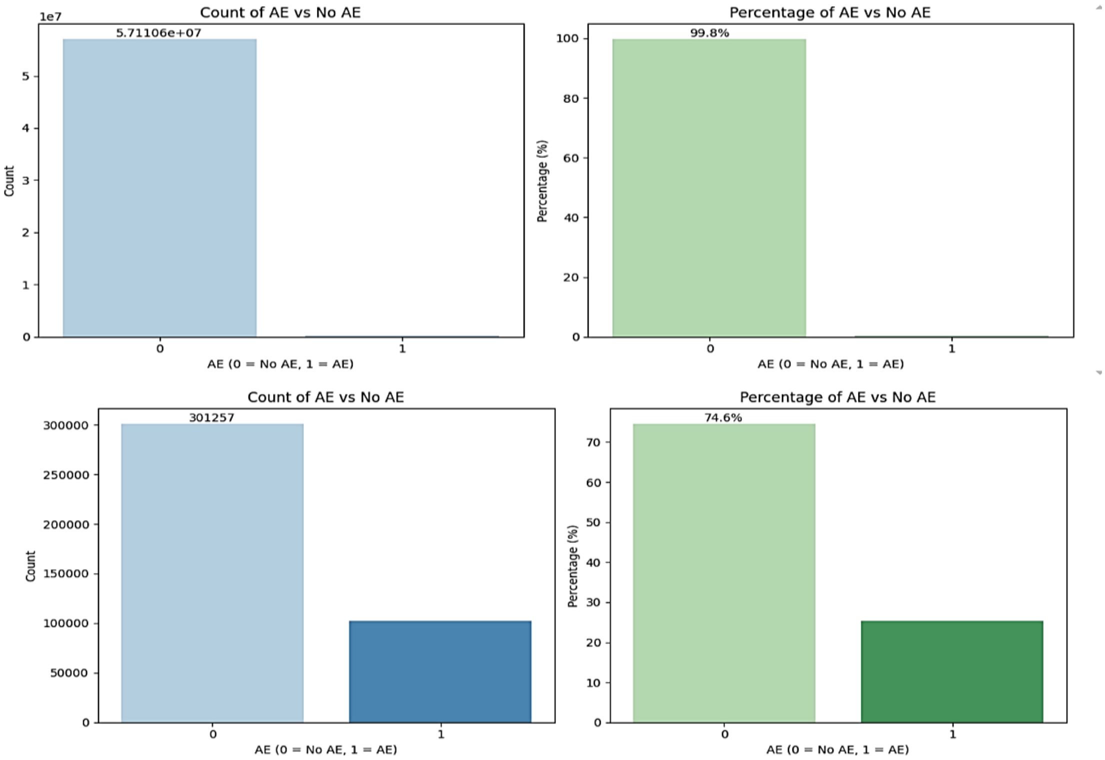

Figure 4. Balancing dataset. Top: before balancing, Bottom: after balancing.

Subsequently, a custom 20-minute binning scheme was applied to group data points into discrete time intervals aligned to 0, 20, or 40 minutes past the hour. This binning facilitates feature aggregation and reduces temporal noise.

Finally, the processed data was aggregated by pot number and 20-minute time groups using a set of descriptive statistics tailored for each variable. Minimum, maximum, and mean functions were applied selectively to capture the range and central tendency of the sensor readings as shown in Figure 5.

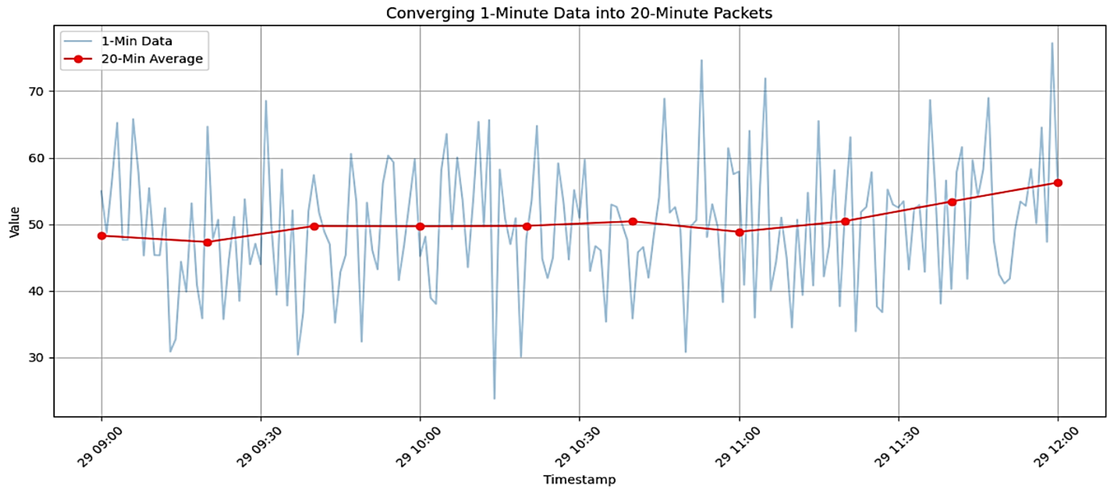

Figure 5. Convergence of high-frequency data into 20-minute time packets.

This preprocessing pipeline ensures that the resulting dataset is clean, temporally aligned, and enriched with meaningful features, enabling robust modelling of AE phenomena in aluminium electrolysis

4. Model Types Evaluated

Several machine learning models were evaluated to determine the most effective approach for predicting anode effects. Gradient boosting algorithms, specifically XGBoost and LightGBM emerged as the top performers. These models demonstrated a strong balance between precision and recall while maintaining computational efficiency, making them the preferred choice for deployment.

Random Forest was also tested and showed a high recall, indicating its effectiveness at capturing true positives. However, this came with reduced precision, leading to a higher number of false positives, which made it less ideal for the application.

Logistic Regression was included as a baseline model. While it was straightforward to implement and provided interpretability, it did not achieve the desired performance levels when compared to the more advanced ensemble methods.

5. Model Performance

The Anode Effect (AE) predictive model was developed and evaluated using a supervised learning approach. The data was split into a training and testing sets in a 70:30 ratio, with cross-validation employed to ensure the robustness and generalizability of the model. The model selected for this task was XGBoost, a high-performance gradient boosting algorithm, which was tuned to operate at a classification threshold of 0.36 to optimize for early detection.

Table 1. Model performance comparisons.

Metric	Random Forest (Default)	Random Forest (Threshold = 0.36)	LightGBM	XGBoost (Threshold = 0.36)
Accuracy	0.8528	0.8336	0.8691	0.854
ROC AUC Score	0.9243	0.9243	0.9428	0.936
Precision (Class 0)	0.85	0.93	0.88	0.94
Recall (Class 0)	0.97	0.84	0.96	0.86
F1-score (Class 0)	0.91	0.88	0.92	0.90
Precision (Class 1)	0.85	0.63	0.84	0.67
Recall (Class 1)	0.51	0.81	0.60	0.83
F1-score (Class 1)	0.64	0.71	0.70	0.74
True Negatives (TN)	58 331	50 602	57 878	51 893
False Positives (FP)	1881	9610	2334	8 319
False Negatives (FN)	9999	3816	8230	3 462
True Positives (TP)	10 484	16 667	12 253	17 021

Performance-wise, the model presented a good trade-off between sensitivity and precision. For Class 0 (No AE), the precision is 0.94, recall is 0.87, and the F1 – score is 0.90, indicating that it is reliable for the identification of normal operational conditions. For Class 1 (AE Likely), the model achieved a precision of 0.67, a recall of 0.83, and thus an F1-score of 0.74 as shown in Table 1. Together, these metrics indicate that the model can detect Anode Effects while keeping a reasonably low rate of false positives.

The model had an overall accuracy of 85.4 %, meaning that it predicted correct operational state in more than 85 out of every 100 cases. The test of the model's ability to discriminate between the two classes was done by the ROC – AUC (Receiver Operating Characteristic - Area Under the Curve) score, with a score of 0.936, which communicates an excellent classification performance as shown in Figure 6.

The confusion matrix says it all:

- 51 893 true negatives (No AE correctly predicted),
- 8 319 false positives (predicted AE when there was not any),
- 3 462 false negatives (not recording AEs), and
- 17 021 true positives (correctly predicted AE).

These results confirm that the model is well-calibrated and highly effective for early prediction of Anode Effects in industrial settings.

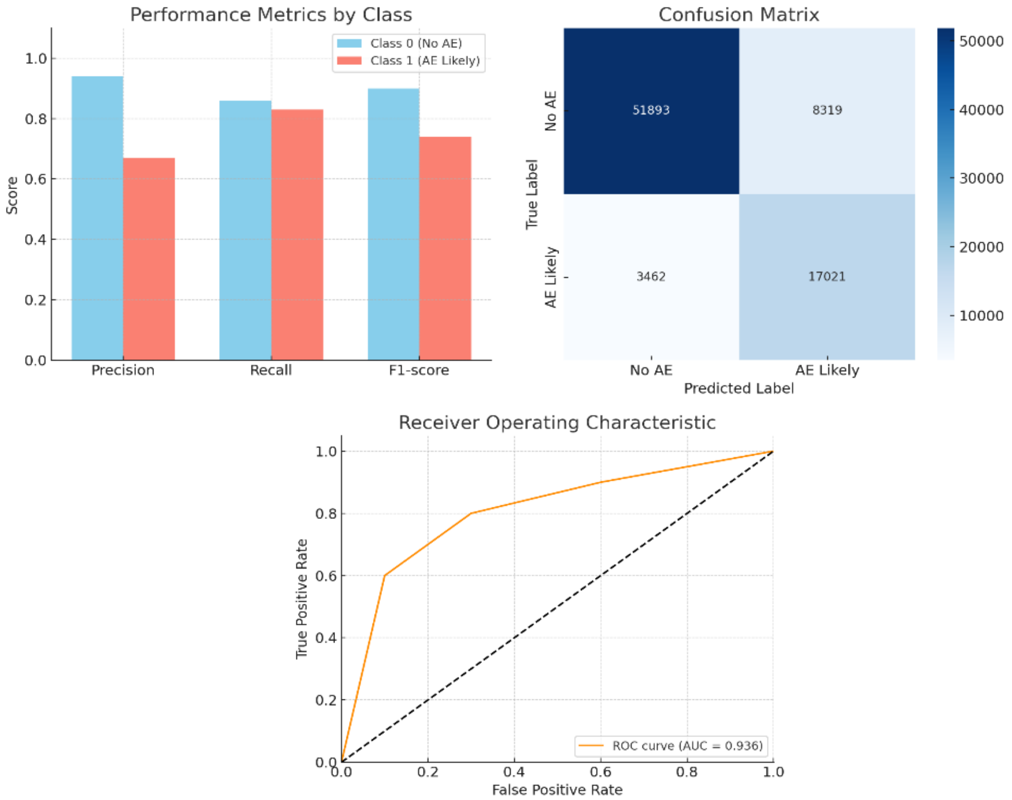

Figure 6. Model performance. Top left: performance metrics by class, Top right: confusion matrix, Bottom: receiver operating characteristics.

6. Real-Time Deployment

The trained model has been integrated into a real-time software system developed in-house. Every 20 minutes, the system scans each pot for (AE) risk. When the model detects a high AE risk, a warning is triggered and immediately displayed in the control room. In addition to this alert, the system also produces recommended operational actions for the plant operators to initiate in order to mitigate the risk pre-emptively, such as initiating alumina feeding (Figure 7). Retraining occurs weekly to uphold the accuracy of the model by ensuring its responsiveness to the fluctuations in plant conditions and operational patterns. Thus, retraining allows the model to adapt to data drift and evolving dynamics in the potline environment.

Our Anode Effect prediction model does trigger some false alarms, about one in every three alerts. However, in real plant operations, these early warnings usually help more than they hinder. When operators respond to an alert, it often stabilizes the pot and can even prevent a real AE from occurring later. So, even if an alarm turns out to be a false positive, it still benefits the process. We built and refined the model using XGBoost to detect AEs early, which is quite reliable, about 85 % accurate, with a strong track record of identifying real events without overwhelming the team with excessive warnings. Every 20 minutes, it checks all pots and provides useful suggestions. We also retrain it weekly to ensure it stays current with ongoing changes in the plant.

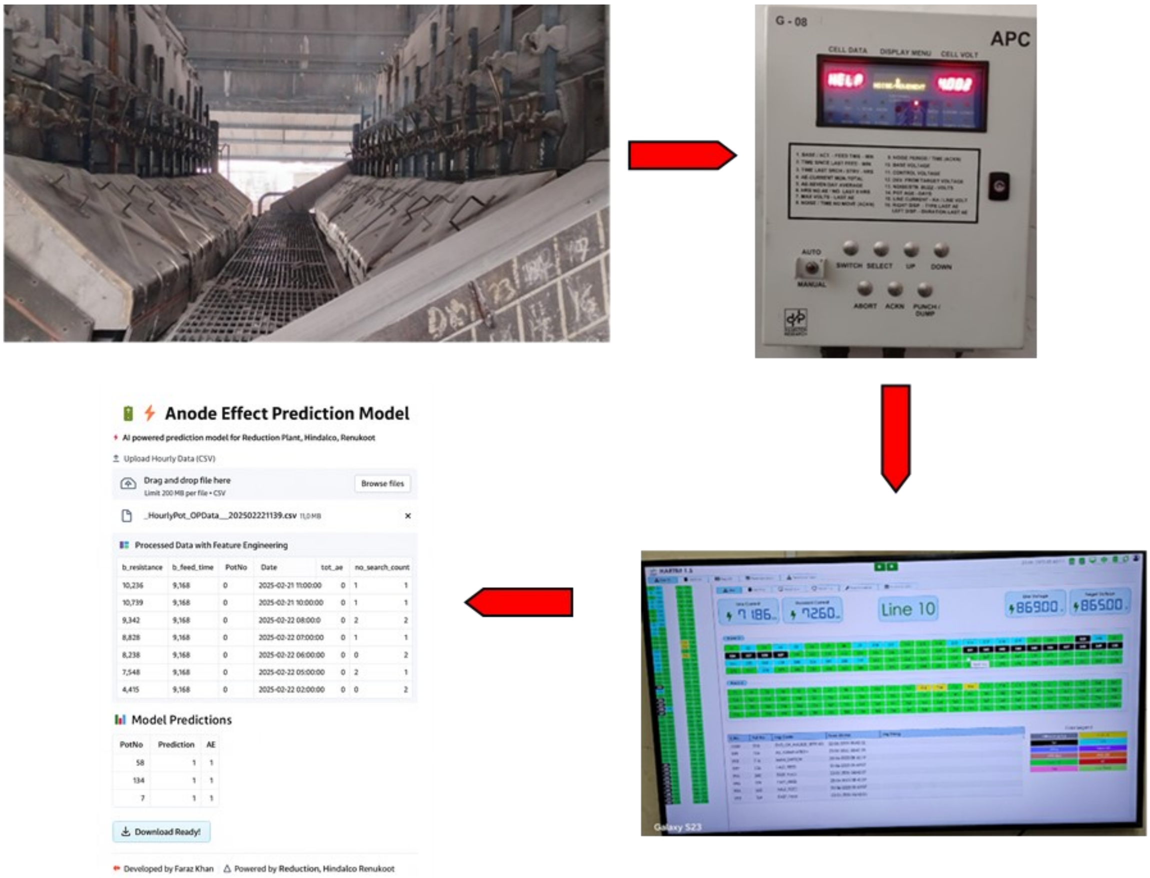

Figure 7. Data flow from pot to APC and from APC to modelling software.

7. Benefits and Impact

Superimposing the predictive AE model onto the operations has led to significant operational improvements. A lesser incidence of Anode Effects is therefore competitive because it allows potline performance to be maintained in a more stable and efficient manner. The decrease in AE per day in Potline is shown in Figure 8. This will lead to better energy efficiency and better use of electrical energy in this electrolysis process. Therefore, energy consumption per tonne of aluminium is reduced, and in this way energy savings and a reduction in operating costs are enormous.

Most importantly a massive decrease in perfluorocarbon emissions was recorded, in keeping with environmental compliance and sustainability programs. Predictions have improved as pot lives were extended, lessening dependency on manual recoveries, which in turn resulted in decreased maintenance and operating downtimes. Further, this model has also enabled operators through APC server system connectivity to acquire real-time information and predictive alerts for quickly and well-timed decision-making on the site. The AE generally results in high temperature of pot, energy loss and PFC emission etc. which can be significantly addressed through this AI-based AE prediction logic (Figure 9).

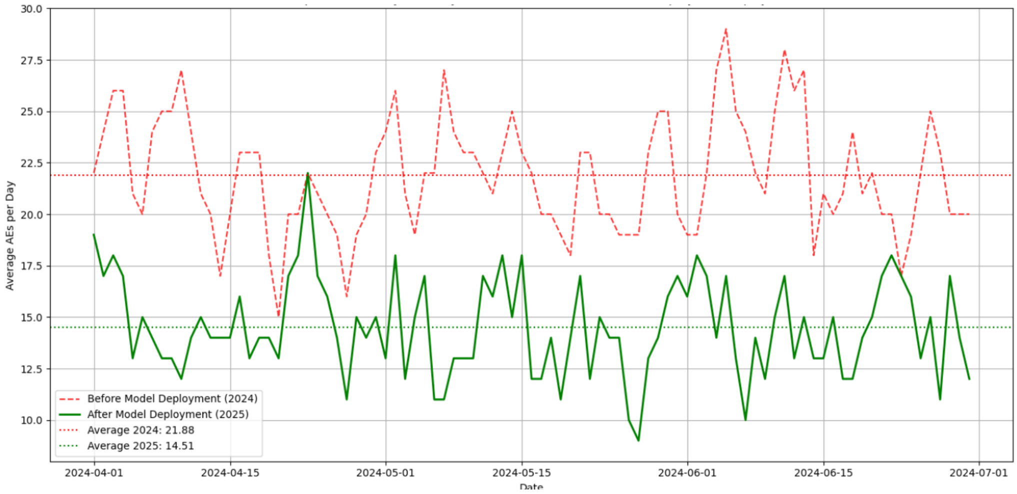

Figure 8. Comparison of daily anode effect numbers before (red dash curve) and after (green curve) model deployment for one potline April-June for 2004 vs 2005.

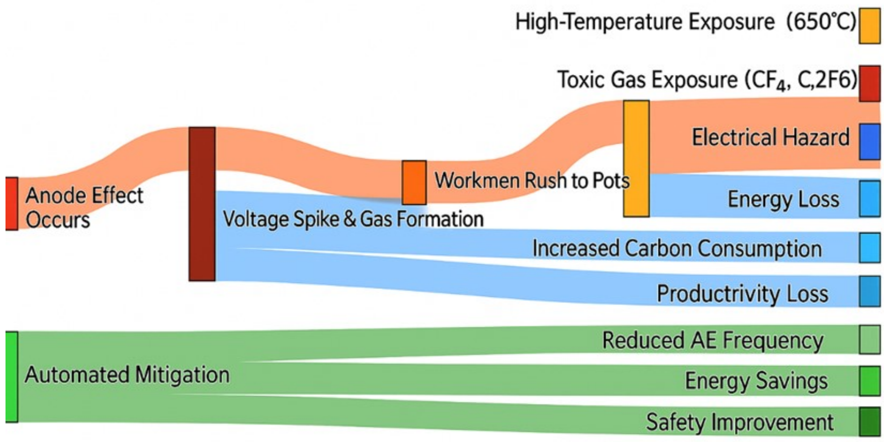

Figure 9. Diagram for anode effect and AI-based mitigation.

8. Conclusions

Within the Renukoot Smelter, the installation of a machine learning-based Anode Effect (AE) prediction model represents a giant leap in the application of AI for intelligent plant functioning. This novel model is unique upon its construction and custom-built to address the specific needs and challenges posed by potline performance monitoring in such an industrially challenging environment.

What makes it even more valuable is the ability to give operators a reliable 20-minute advanced warning of the likelihood of anode effect occurrence. This gives the team ample time to intervene and carry out preventive work to minimize disruption, energy consumption, and greenhouse gas emissions. As a result, the plant can operate more efficiently while offering tighter control of costs and a smaller environmental footprint.

The model is designed for continuous learning and improvement by retraining itself with the latest process data. This enables adaptation to any changes in the smelter operations and maintains accuracy over time. Operator and expert feedback are also incorporated into the process, improving the reliability and responsiveness of the system.

This AI – led solution boldly heralds a futuristic approach to industrial automation-and it is helping to set a precedent as to what predictive technologies can do in the aluminium world.

9. References

1. George E. Totten et al., *Handbook of aluminum: Volume 2 alloy production and materials manufacturing*, 2003, 736. https://doi.org/10.1201/9780429223259
2. Geir Martin Haarberg, Effects of electrolyte impurities on the current efficiency during aluminium electrolysis, *Proceedings of 33rd International ICSOBA Conference*, Dubai 2015, Paper AL27.
3. Warren Haupin and Edward J. Seger, Aiming for zero anode effects, *Essential Readings in Light Metals* 2016, 767–773. https://doi.org/10.1007/978-3-319-48156-2_114
4. Kai Grjotheim and Halvor Kvande, *Introduction to aluminium electrolysis*, DIN Media GmbH, 1993.
5. J. Thonstad, T.A. Utigard and H. Vogt, On the anode effect in aluminum electrolysis, *Essential Readings in Light Metals* 2016, 131–138. https://doi.org/10.1007/978-3-319-48156-2_18
6. David S. Wong, Alton T. Tabereaux and Pascal M. Lavoie, Anode effect phenomena during conventional AEs, low voltage propagating AEs & non-propagating AEs, *Light Metals* 2014, 529–534. https://doi.org/10.1007/978-3-319-48144-9_90
7. A. Mulder, P. Lavoie and R. Düssel, Effect of early abnormality detection and feed control improvement on reduction cells performance, *Proceedings of the 12th Australasian Aluminium Smelting Technology Conference*, Queenstown, New Zealand, 2018, 2–7 December 2018.
8. Andri Ísak Þórhallsson, Anode effect reduction at Nordural — practical points, *Light Metals* 2015, 539–543. https://doi.org/10.1007/978-3-319-48248-4_90
9. Sylvain Fardeau et al., Statistical evaluation and modeling of the link between anode effects and bath height, and implications for the ALPSYS pot control system, *Light Metals* 2014, 845–850. https://doi.org/10.1002/9781118888438.ch142
10. Ron Kremser et al., Investigation of different measurement techniques for individual anode currents in Hall-Héroult cells, *Proceedings of the 12th Australasian Aluminium Smelting Technology Conference*, Queenstown, New Zealand, 2018, 2–7 December 2018, 2–7 December 2018.
11. Zuguo Chen et al., Anode effect prediction based on collaborative two-dimensional forecast model in aluminum electrolysis production, *Journal of Industrial & Management Optimization* 2019, 15(2), 595–618. https://doi.org/10.3934/jimo.2018060
12. Lukas Dion et al., Preventive treatment of anode effects using on-line individual anode current monitoring, *Light Metals* 2017, 509–517. https://doi.org/10.1007/978-3-319-51541-0_64
13. C. Cheung et al., Frequency response analysis of anode current signals as a diagnostic aid for detecting approaching anode effects in aluminum smelting cells, *Light Metals* 2003, 887–892. https://doi.org/10.1002/9781118663189.ch151
14. Fernando Pestana da Costa et al., Computer algorithm to predict anode effect events, *Light Metals* 2012, 655–656. https://doi.org/10.1007/978-3-319-48179-1_112
15. A. Wilson, M. Illingworth and M. Pearman, Anode effect prediction and pre-emptive treatment at Pacific Aluminium, *Proceedings of the 12th Australasian Aluminium Smelting Technology Conference*, Queenstown, New Zealand, 2018, 2–7 December 2018.
16. Yu Zhang, Study on anode effect prediction of aluminium reduction applying wavelet packet transform, *Proceedings of 6th International Conference on Intelligent Computing*, China, August, 2010, August 18-21, 2010, 477–484. https://doi.org/10.1007/978-3-642-14831-6_62

17. Nazatul Abd Majid, *Cascade fault detection and diagnosis for the aluminium smelting process using multivariate statistical techniques*, PhD Thesis, The University of Auckland, 2011. http://api.digitalnz.org/records/30143110/source
18. Zhixin Zhang et al., Anode effect prediction based on expectation maximization and XGBoost model, *Proceedings of the 2018 IEEE 7th Data Driven Control and Learning Systems Conference*, Enshi, China, 2018, 25–27 May 2018, 560–564. https://doi.org/10.1109/ddcls.2018.8516046
19. Ron Kremser et al., Anode effect prediction in Hall-Héroult cells using time series characteristics, *Applied Sciences* 2020, 10, 9050. https://doi.org/10.3390/app10249050