An Ensemble Soft Sensor for SiO_2 Grade Prediction in Iron Ore Reverse Cationic Flotation

A T+30 minute concentrate-grade prediction architecture stacking LSTM, 1D-CNN, and Random Forest models on physics-derived, reagent-intensity, and temporal-lag features, producing an explainable signal for proactive reagent control.

AUTHORS	AFFILIATION	CLASSIFICATION
T. Davhana	**Axiom Process Technology**	**Public Release**
D. Mthimunye	**Nelspruit, South Africa**	**Vol 01, No 02**

ABSTRACT

Reverse cationic flotation is the dominant route for upgrading iron ore concentrates, with the silica content of the concentrate determining the saleable grade and the corresponding price realisation. The principal closed-loop measurement of SiO_2 grade is laboratory X-Ray Fluorescence assay on grab samples, with turnaround times of two to four hours from sampling to availability. The control implication is that all reagent adjustments are retrospective by definition. This paper presents an ensemble soft sensor architecture that predicts SiO_2 grade thirty minutes into the future from a continuously available SCADA feature vector. The architecture stacks an LSTM network for long-range reagent-residence memory, a one-dimensional Convolutional Neural Network for local temporal patterns, and a Random Forest regressor for feature-importance transparency. A ridge meta-learner produces a point estimate together with a ninety-five percent prediction interval. Approximately forty-five to sixty-five engineered features are organised into five categories: physics-derived flotation kinetics, reagent-intensity normalisations, feed characterisation, temporal lag windows reflecting reagent residence time, and statistical-window stability indicators. The sensor enables proactive amine, starch, wash-water, and froth-depth adjustments *before* the laboratory assay confirms the disturbance.

KEYWORDS

Soft sensor, iron ore, reverse cationic flotation, silica, LSTM, 1D CNN, Random Forest, ensemble stacking, prediction interval, proactive control.

1. Introduction

Iron ore reverse cationic flotation upgrades the concentrate by floating quartz and silicate gangue with an amine collector while depressing iron-bearing minerals with starch. The grade-recovery placement of the operating point is sensitive to the air flow per cell, froth depth, the dosage and ratio of amine to starch, the feed pH, and the hydrodynamic regime in the collection and froth zones. Each of these variables is continuously instrumented at the DCS level, yet the only direct measurement of the objective function, the silica content of the concentrate, arrives by laboratory XRF assay on the order of hours after the sample is drawn. The result is that the operator is steering a process whose principal output is observed retrospectively.

The cost of this latency is twofold. Grade departures that exceed contract specification attract penalties or rejection at the customer; grade overshoots above specification represent recovery sacrificed to a conservatism the operator has no faster way to unwind. The economic prize for a thirty-minute forward estimate of concentrate SiO_2 grade is the elimination of both ends of this distribution. A soft sensor that closes this gap must be continuously available, physically interpretable to the resident metallurgist, and explainable to the operator who acts on it.

2. Process and Data Architecture

2.1 SCADA input set

The architecture is fed by twelve continuously available SCADA tags at a one-second DCS scan rate. Per cell, the air flow rate is supplied by a rotameter or mass-flow transmitter; the froth depth or column level by an ultrasonic or float sensor. Feed flow rate is taken from a magnetic flow meter, feed density and percent solids from a nuclear density gauge, and feed pH from an inline electrode. Amine collector and starch depressant dosages are read from dosing-pump flow transmitters. Wash-water flow per column, concentrate flow, tailings flow, and slurry temperature complete the measurement set. Feed SiO_2 and iron grade from periodic XRF assays serve as training labels only and do not appear as real-time inputs.

2.2 Preprocessing

Raw SCADA streams are subjected to timestamp alignment, interquartile-range outlier removal at three sigma, forward-fill on laboratory channels to maintain a continuous training label, and min-max normalisation. The train, validation, and test split is strictly chronological. The test set is constructed to deliberately include ore transitions, ensuring that generalisation is evaluated against the regime changes the sensor must survive in production.

3. Feature Engineering

Five feature categories combine to a vector of approximately forty-five to sixty-five features, presented in parallel to each base learner.

3.1 Physics-derived features

Superficial gas velocity is computed as J_g = Q_air ÷ A_cell, governing bubble generation and gas holdup in the collection zone. Froth residence time is computed as τ_f = h_froth ÷ v_froth, governing drainage and entrainment behaviour in the froth phase. Hydraulic bias is computed as B = Q_wash –

Q_conc; a positive bias drives net downflow and is associated with cleaner iron concentrate.

3.2 Reagent-intensity features

Specific amine dosage D_a = Q_amine ÷ Q_feed and specific starch dosage D_s = Q_starch ÷ Q_feed normalise reagent consumption to feed rate in grams per ton. The selectivity index SI = D_a ÷ D_s captures the collector-depressant balance as a single dimensionless ratio and serves as a direct proxy for silica selectivity.

3.3 Feed characterisation features

Feed solids content is derived from the nuclear density gauge through the standard volumetric relation. Volumetric throughput Q_v = Q_feed ÷ ρ_slurry captures the actual volume flow into the cell and is the correct turbulence proxy in the collection zone. A pH-deviation feature ΔpH = pH – pH_target encodes the distance from the amine optimum, conventionally near pH 10.5, where selectivity is maximal.

3.4 Temporal lag features

Reagent residence time in reverse cationic flotation is approximately twenty to thirty minutes. The feature set therefore includes lagged values of air flow at five, ten, and twenty minutes; lagged amine dosage at ten, twenty, and thirty minutes; lagged froth depth at five, ten, and twenty minutes; and lagged feed density at ten, twenty, and thirty minutes. These lags provide the architecture with the memory it needs to attribute present grade outcomes to upstream reagent and ore changes.

3.5 Statistical window features

Fifteen-minute rolling means establish a trend baseline; fifteen-minute rolling standard deviations indicate process stability; an exponentially weighted moving average with λ = 0.1 supplies an exponentially decaying memory for slow drifts. Rate-of-change features ΔX ÷ Δt over a five-minute horizon function as leading indicators of ore transition.

4. Model Architecture

4.1 Base learners, parallel training

Three base learners are trained in parallel on the full feature vector. The LSTM network operates on sequences of thirty time steps and is intended to capture the long-range reagent-residence memory that the lag features expose. The 1D Convolutional Neural Network uses kernels of five to fifteen steps and is intended to identify local temporal patterns in the SCADA streams. The Random Forest regressor uses five hundred trees and contributes feature-importance ranking that makes the contribution of each engineered feature transparent to the metallurgist.

4.2 Ensemble stacking layer

The three base-learner outputs are combined through a ridge meta-learner that produces both a point estimate and a ninety-five percent prediction interval. The ridge weights are updated on a rolling validation window so that the relative confidence in each base learner adapts to the regime the plant is currently in. Where the LSTM and 1D CNN disagree sharply, the prediction interval widens; the metallurgist is alerted that the sensor itself is uncertain, not given a single number with no context.

4.3 Output and control integration

The output is the predicted concentrate SiO_2 grade at T+30 minutes, updated every DCS scan, accompanied by a ninety-five percent prediction interval and the SHAP-style attribution from the Random Forest leg of the ensemble. The control layer above the sensor responds proactively: amine dosage is adjusted up or down, starch dosage is rebalanced, wash water is trimmed, and the froth-depth setpoint is shifted before the laboratory assay confirms the disturbance. The lab assay returns as a delayed training label, closing the supervisory loop, not the regulatory one.

5. Architectural Position in the Axiom Control Stack

The SiO_2 soft sensor occupies the third layer of the Axiom control hierarchy, alongside the ore-hardness soft sensor and other physics-informed estimators. It sits above a healthy DCS and above the anomaly-detection layer that protects the regulatory foundation from sensor dropout, sticky valves, and dead-banded transmitters. The sensor output enters the Advanced Process Control layer as a feedforward signal, where the hybrid FLC and MPC stack converts predicted grade departures into reagent and air-flow moves. The AI Optimisation Layer above the APC consumes both the prediction and the prediction interval when computing the thirty-minute economic setpoint recommendations that the metallurgist accepts or rejects.

6. Conclusion

The laboratory turnaround on concentrate SiO_2 assay is the binding latency in iron ore reverse flotation control. The ensemble soft sensor presented in this paper closes that gap by combining physics-derived, reagent-intensity, feed-characterisation, temporal-lag, and statistical-window features through a stacked LSTM, 1D CNN, and Random Forest architecture. The output is a continuously available, prediction-interval-qualified, feature-attributable forecast of concentrate grade thirty minutes ahead. The control consequence is the elimination of both grade penalties and recovery sacrifices attributable to the assay latency.

References

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), pp. 1735–1780.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), pp. 5–32.

Lundberg, S.M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, pp. 4765–4774.

Wills, B.A. and Finch, J.A. (2016). *Wills' Mineral Processing Technology: An Introduction to the Practical Aspects of Ore Treatment and Mineral Recovery*, 8th ed. Butterworth-Heinemann, Oxford.