Machine Learning for Blast Furnace Productivity Improvement at Jindal Steel and Power Angul

Digital technologies are transforming industry at all levels. Steel has the opportunity to lead all heavy industries as an early adopter of specific digital technologies to improve our sustainability and competitiveness. This column is part of AIST's strategy to become the epicenter for steel's digital transformation, by providing a variety of platforms to showcase and disseminate Industry 4.0 knowledge specific for steel manufacturing, from big-picture concepts to specific processes.

Authors

Anil Gandhi
Data Scientist and President, Qualicent Analytics Inc., Santa Clara, Calif., USA
anil.gandhi@qualicent.net

Ritesh Mishra
Jindal Steel and Power Ltd., Angul, OD, India
rmishra@angul.jspl.com

Aniruddha Fadnis
Qualicent Analytics Inc., Santa Clara, Calif., USA

Richa Jaiswal
Jindal Steel and Power Ltd., Angul, OD, India

Sambit Mishra
Jindal Steel and Power Ltd., Angul, OD, India

Amitava Sircar
ABS Technical Services, Bokaro, JH, India
amitava.sircar@gmail.com

Jindal Steel and Power Ltd. (JSPL), a major steel producer in India, commissioned its first blast furnace, with a size of 4,554 m^3, in May 2017, at its plant located at Angul, Odisha, India. This paper will describe machine learning (ML)–based computational efforts for improved blast furnace productivity, complementing the knowledge of iron and steel specialists. A productivity trend for the evaluation period is shown in Fig. 1. Capacity of the plant went from 1.5 MTPA (with a production route utilizing direct reduced iron (DRI), an electric arc furnace (EAF), continuous caster (CC), rolling mill) in 2017 to 6 MTPA (utilizing a blast furnace (BF), basic oxygen furnace (BOF), CC, rolling mill production route) at end of the evaluation period.

BF processes are influenced by many factors, with quite a few of them interacting with each other. The process is immensely complex. Such process complexity in modern state-of-the-art blast furnaces enables higher quality of hot metal; however, this complexity can also result in patterns of BF behavior that become difficult to capture by conventional means. In turn, this can lead to lower productivity at times. It is in this context of hidden, intractable and nuanced sources of furnace behavior and in the context of process complexity that machine-learning methods become a useful device and a vehicle for information discovery. Fortunately, Industrial Internet of Things (IIoT) has enabled modern blast furnaces with a large number of continuous data streams that can be harnessed and exploited to great benefit, as this paper will show.

Furnace behavior has been examined in the past with various statistical multi-variate (MVA) methods by practitioners in the steel industry. These statistical methods have been around for quite some time and to a reasonable degree they have provided benefit to the consumer of this information. The theory and application for design of experiments (DOE), standard regression and other statistical methods are well developed, and it would be superfluous to address it here. However, there are limitations to these traditional methods when (1) the number of features (variables) is large, (2) many of these features are correlated, (3) there are deep interactions among many features and (4) the number of data points is small relative to the number of features. In these limiting conditions, ML methods play an important role in improving productivity and quality in steel manufacturing.

Although this paper will refer to the specifics of BF processes, methods and algorithms employed in this analysis should be applicable throughout the steel plant. And while data science concepts are used quite extensively in this analysis, graphical devices are deliberately relied on to highlight findings, possibly at the expense of some mathematical rigor, but of more use to the practitioner in the steel industry. Machine-learning methods have come of age and it is believed that ML, together with more sophisticated methods such as deep learning, will be employed routinely in all aspects of steel manufacturing, leading to simultaneous increases in productivity and quality.

Summary

From the models utilized, it was found that permeability (K) of the blast furnace plays a key role. Of course, the importance of K has been known since the advent of blast furnaces; however, since there are no direct measurements possible inside the column of a blast furnace, quantification has been lacking. Instead, historically, practitioners have resorted to indirect extrapolation from other measures to compute the implied value of K. Machine learning allows us to discover all the complex associations of K to a finer degree than was otherwise and previously possible.

Permeability as Proxy Variable — Empirical evidence has been found that not only does permeability K play a pivotal role in influencing productivity, but in fact it provides a pathway for many other variables to influence the process. In other words, K acts like a proxy variable for productivity. This empirical finding, namely K acting as a modulating variable, is not entirely surprising, given that movement of gases in the furnace has been implicated in various physics-based models as the key factor, and in fact, has been the subject of much modeling with computational fluid dynamics (CFD) and other methods.[1]

Alternative Model Representations — At the outset, the authors note and distinguish between non-parametric empirical discovery of effects (through ML in this case) and physics-based models. In this paper, the focus is on non-parametric empirical discovery (also called model-less representation). All discoveries must eventually be explained through physics for root cause and corrective action and improvement; however, purely for predictive purposes, it is not necessary to know the physics but only that the model has high confidence and is predictive enough. This is a philosophical point, and outside the scope of this paper; however, it is mentioned only to bring to the attention of the reader that these philosophical renderings can influence the interpretation of results. The term physics is used as a comprehensive term that includes physics, metallurgy, chemistry, mechanical and any other related discipline that describes phenomena in a parametric model.

Variables of Importance — Among the variables of considerable influence are coke strength, ash and moisture. Less expectedly, the amount of lime in sinter, as part of the charge materials mix, has been found to play a key role in productivity. It is reasonable to assume that since iron is the payload in sinter as the element of interest, a higher percentage of iron relative to other components would be desirable. Instead, it has been found that an optimum amount of lime must be present for higher productivity; this comes at the expense of reduced fraction of iron, which is counterintuitive. This does not mean that higher lime percent is a desirable feature but rather the data (restrictively) suggests only that an optimum level of lime percent is necessary for a given quality of sinter to extract higher productivity. Alternatively, there must be balance in sinter chemistry (fluxed enough for meeting chemical requirements and physical strength). ML allowed for the desired CaO addition level required for different kinds of raw material mix in a sinter plant to be extracted.

It is believed that this phenomenon results from the binding properties of lime,[2] which increases the strength of sinter and prevents mechanical degradation in the blast furnace. Mechanical degradation results in smaller particles (fines) which effectively block the free movement of gas in the furnace. All of

Figure 1

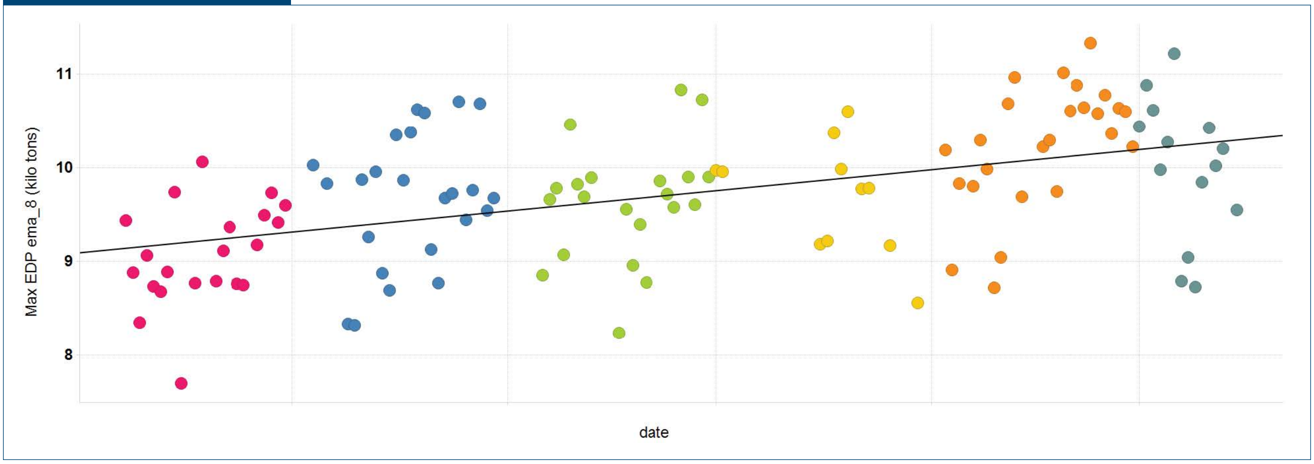

Productivity trend in the evaluation period for the blast furnace at the Jindal Steel and Power Ltd. plant at Angul, Odisha, India.

this relates to K, the permeability metric, confirming once more the pivotal role that permeability plays in improving productivity. The optimum percentage of lime depends on the quality of sinter, so that higher-quality sinter, which is less prone to mechanical degradation, requires less lime percent and therefore delivers higher productivity, but, for a given level of sinter mechanical quality, an optimum percentage of lime is required for higher productivity. More on this in the Results and Discussion section.

The top few variables of importance (varimps) are shown in Fig. 2 and the role of these varimps and other variables emerging out of the ML model will be discussed. The table suggests that impairment in one of these inputs due to incoming quality can be compensated for by better quality in the other metrics to some extent — if the total permeability can effectively be maintained at reasonably good levels.

Data Set Description

Data set measurements were recorded in a stable evaluation period, where there were no major unscheduled or scheduled downtimes. Five different kinds of data sets are extracted and processed — tap/cast, charge, quality, L2 and burden distribution.

Data Aggregation — Hot metal output in the evaluation period is computed for every hour and recorded as equivalent daily production (EDP) in tons/day, which equates to production for each hour multiplied by 24. This is shown in Fig. 3; the left figure showing a snippet of time trend along with one-sided exponentially weighted moving average, the figure on the right showing the density distribution. ma_0 refers to the raw quantity of production for the specific hour; clearly this is a very noisy signal — some of the noise is due to a lapse in accurate production accounting for each tap event. Each individual tap or casting output can be undercounted or overcounted to some degree, but when combined over several taps or combined over several hours, the accounting errors disappear because the undercounting or overcounting output in each tap event must be accounted for at some point. ema_8 in the figure on the right represents the mean of hot metal output exponentially weighted 8 hours in each direction while the left figure shows a one-sided ema_8.

Model accuracy and predictiveness depend on capturing most of the temporal signal and discarding most of the noise. Increasing averaging length decreases the signal by decreasing the captured granularity of temporal variation but it also reduces the noise coming from the raw signal. Modeling was performed with different levels of EDP averaging length and two-sided 8 hours ema was found to offer the best model accuracy. Consistent with that, the figure on the right in Fig. 3 visually suggests that noise from accounting is largely eliminated with a time averaging length of 8 hours with no fat tails remaining. The distribution shrinks at a much slower rate beyond this. Therefore ema_8 is chosen as the "target" variable (the variable to maximize).

Output by Taphole — Various hot metal output metrics are shown in Fig. 4 as part of exploratory analysis to understand the bearings of the data set, which is usually the first step before any ML exercise is undertaken. Fig. 4 defines the characteristics of this furnace by taphole; every furnace will have its own characteristics.

Raw materials (charges) are continuously fed into the furnace and its measurement is done on the weigh

Figure 2

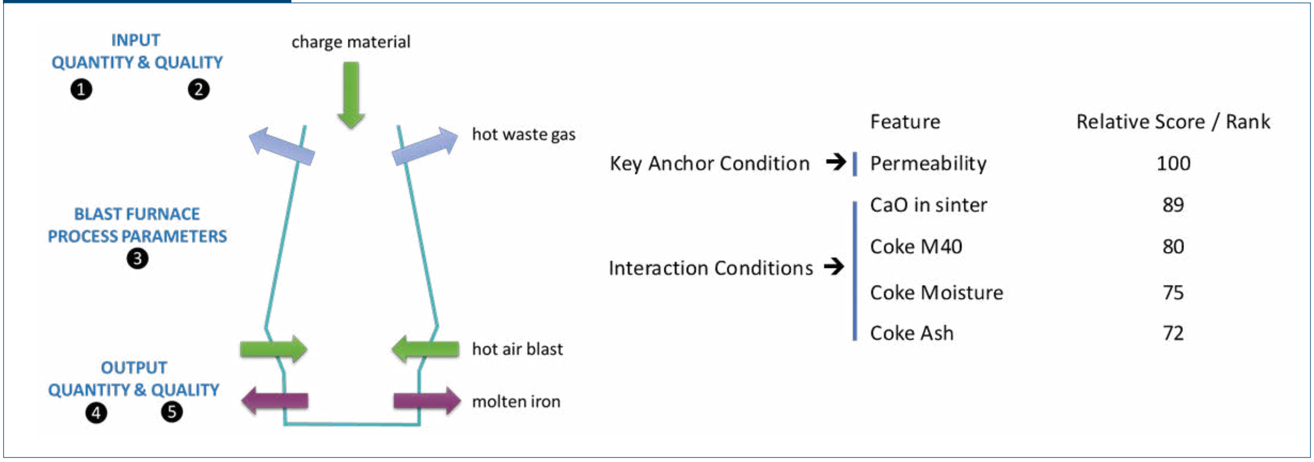

Top five key features (variables) in the machine-learning (ML) model that describe productivity of the blast furnace.

Figure 3

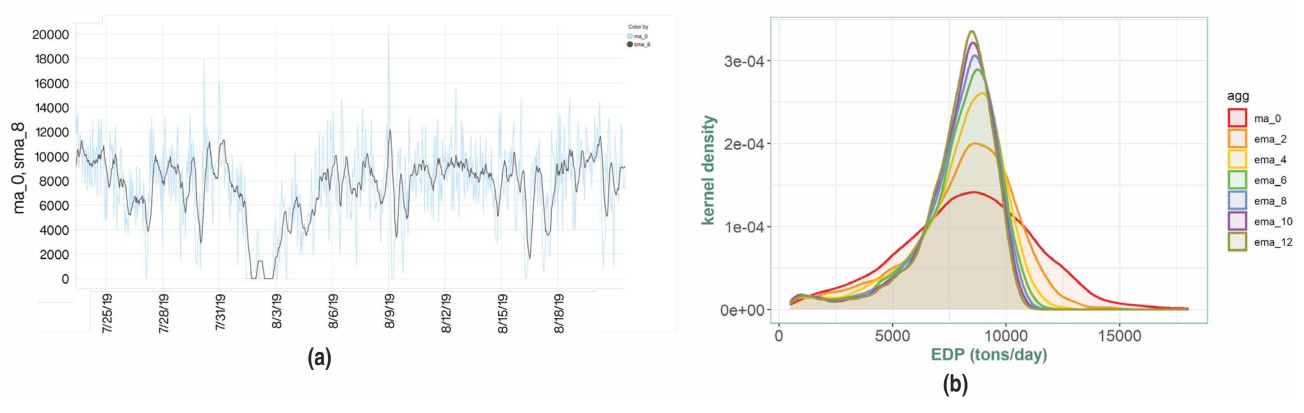

Equivalent daily production (EDP) by hour time trend – snippet (a) and EDP distribution for various moving averages (b).

hoppers on a pre-set schedule. Some of the raw materials such as sinter and coke are generated on campus in the integrated steel plant at JSPL while some other materials like pellets are imported into the plant. Prime coke consumption (a key indicator of blast furnace economics) was lowered to almost 350 metric tons/day over the 18-month commissioning and there remains room to make further improvements in this metric: all this while hot metal output increased.

Sampling Adequacy — Quality parameters for raw materials are measured at the stockhouse at various intervals. Some parameters are measured up to twice a shift while others may be measured once a day or once every two days. Parameters that have an adequate sampling rate were chosen. The data table constructed at a clocking index of 1 hour records the average value of quality parameters and blast furnace process parameters (as measured by L2) for every hour, while charge data is summed for every hour. An 8-hour shift, therefore, has eight measurement opportunities so that parameters that are measured once a shift will rank 12.5% (=1/8) on the y axis in Fig. 5; likewise, about 4.2% (1/24) on the y axis for measurements that are taken once a day.

Out of 182 quality features, about 41% of them that are sampled at a frequency of once a day or better are selected, as shown in Fig. 5, and for these features, values for the remaining hours are imputed based on "down fill" methodology — i.e., previous values stay valid until the new measurement is made. Thirty-two L2 parameters describing blast furnace measurements were included in the source in the analytics table e.g., top pressure, hot blast temperature, RAFT, permeability, etc.

Figure 4

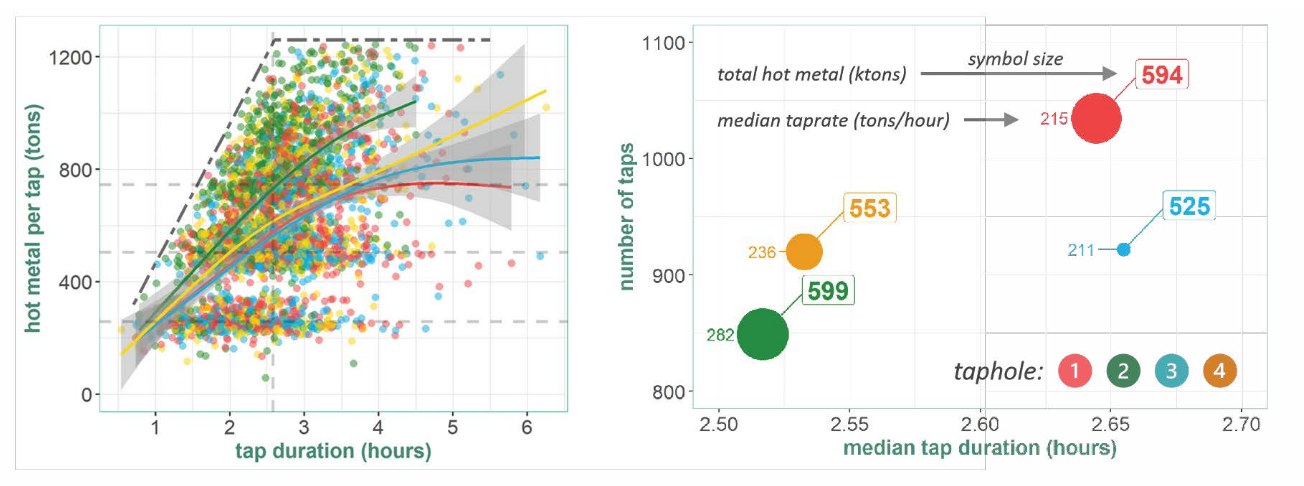

Various hot metal output metrics by taphole, indicating the characteristics of this furnace.

Figure 5

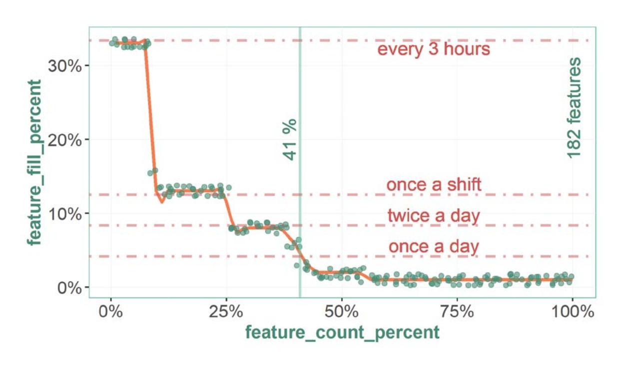

Selection of quality parameters based on sampling adequacy. (Data points have been jittered for visibility.)

Analytics Methodology

Unit of Production — Hot metal generation in the blast furnace is a continuous process; all statistics, therefore, are generally quoted per unit time. Tapping events (reflecting the output variable, also called the supervising variable) are, however, asynchronous; one or more of the four tapholes could be open in each of the hours in the data set. There are two ways in which a raw data set can be prepared for mining (a) each tap event as a unit of production or (b) each hour of production as a unit of production. Both methods were used, but the following discussion refers to the second method because it provided better accuracy.

Lookback Period — To build a mining-ready table, with a target variable (hot metal output) and predictors, each hour must contain the output for that hour and the inputs that supplied into that output hour during the descend and residence time, X, of the materials in the furnace column. It was suspected that X is in the range of 8 hours based on prior domain knowledge of blast furnaces; perhaps not surprisingly it was found X = 8 hours provides the best accuracy in an unbiased ML model, confirming prior belief. It was noted, however, that model accuracy drops off sharply outside ±2 hours of X = 8. The word "unbiased" is used to mean without the assistance of a physics-based model. This, of course, does not mean that there is not a physics-based model (all physical processes do eventually) but only that the discovery methods used here are independent of it.

Analytics Platform and Algorithms — R was used as the platform of choice for algorithmic analysis,[3] although python would be a perfectly good choice as well. Both these platforms are free, and both offer a wide variety of state-of-the-art machine-learning algorithms through packages. Several of these algorithms under the supervised and unsupervised learning paradigm are employed in this analysis to model productivity and permeability of the blast furnace. Xgboost algorithm[4] showed the best promise in overall accuracy and effectiveness, which is not surprising, given the versatility of this algorithm accommodating a multitude of data set conditions including high degree of correlation between many variables. Earth algorithm is also utilized; however, this algorithm requires a thinned data set where the variables have been pared down to a more manageable smaller number of less correlated variables. The output of xgboost is fed to earth algorithm; the advantage in earth is that it allows extraction and visualization of interactions which is one of the objectives in this exercise.

In addition to the above algorithms, other techniques were employed, such as principal components regression, principal component analysis, hierarchical clustering and others. More details on the techniques themselves is widely available on the internet with a simple search.

All data is temporally aligned to each production unit, so that each row in the analytics table, corresponding to an hour of production, has columns that describe the feature values associated with this production unit. All features for this unit of production are synchronized with a lookback period of several hours due to the descend time of material in the blast furnace column that

Figure 6

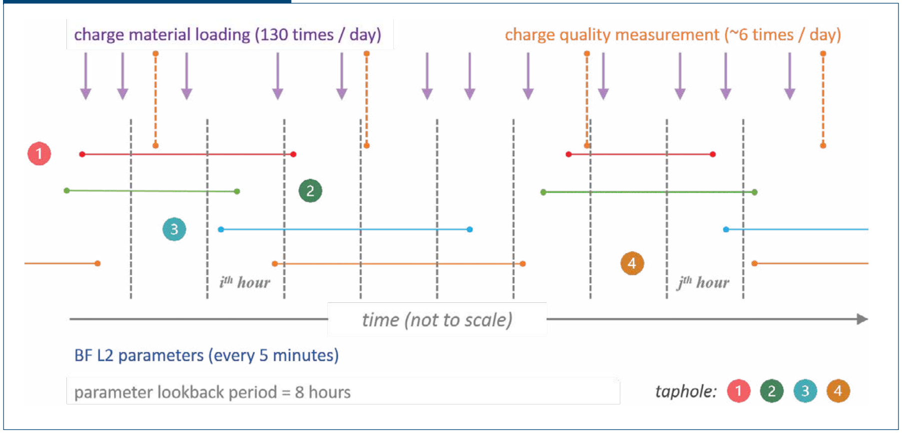

Unitization of hot metal output quantity using each hour as a unit. A lookback period of 8 hours is found optimal.

affect the production and quality in each hour of tap output, as shown in Fig. 6.

Target Variable — The productivity of the blast furnace is defined before it is modeled. For analytics purposes, the authors use the instantaneous value of production, an instant being an hour in this case, since an hour is the smallest aggregated unit of production across all metrics. This hourly production is normalized to a 24-hour period since the plant is used to assessing production in daily units in kilo tons per day. This instantaneous value of production is called equivalent daily production (EDP). The objective for the ML algorithm is to model and maximize EDP or more specifically, the exponential moving average over 8 hours or ema_8_EDP, defined earlier, as the target variable.

Model Overfit Prevention — Next, the issue of regularization is addressed to prevent overfitting. State-of-the-art ML algorithms have the ability to capture every fine nuance in the data set, converging to a perfect fit with the data in the absence of constraints. But this perfectly fit model is also a perfectly bad model. This is so because an unconstrained "perfectly fit" model includes considerable noise, while the intent is to model the signal (hot metal output) and not the noise. Inclusion of too much noise in the model is called overfit. Two mechanisms are used to safeguard against this phenomenon of overfit — test-train separation and n-fold cross-validation.

As a first layer of defense against model overfit, n-fold cross-validation is used. For eight-fold cross-validation, the model is built on seven out of eight folds in the training set and then measured against the remaining eighth fold. This is usually done by the algorithms internally and automatically and done several times repeatedly, measuring against the different folds to prevent overfitting. This prevents random noise from influencing the model. However, signal variation originating from variables that are not explicitly called out in the data set must also be mitigated. And this is done by splitting the data set into a training subset and test subset.

So, the second layer of defense against model overfit is accomplished by splitting the data set into two distinct, time-chunked (i.e., not randomly picked) sub-data sets in a ratio of 70% training subset and 30% test subset, or 80%/20%. The model is built on the training subset with n-fold cross-validation and tested on the remaining members of the data set called the test subset. "Training" is data science speak for developing a model using ML methods with a class of algorithms called supervised learning. Building the model on a chunked training subset and testing on another chunked subset ensures that the model accuracy is tested against unseen data in a different time period when the state of unmeasured variables is potentially quite different.

In addition to these two methods, there are boosting parameters that the user can and should optimize, for example the number of rounds of iteration and the learning rate in each iteration. A discussion on these and other optimizations can be found on the xgboost website.[4]

Results and Discussion

Employing the methodology as described earlier, an ML model is developed for EDP with ~115 predictors or features — the predictors coming from L2, quality and charge data. The model is shown compared against actual data in Fig. 7, achieving a cross-validated and regularized model with r-square of ~0.87 for the test sub-data set. Out of 153 million possible terms,

Figure 7

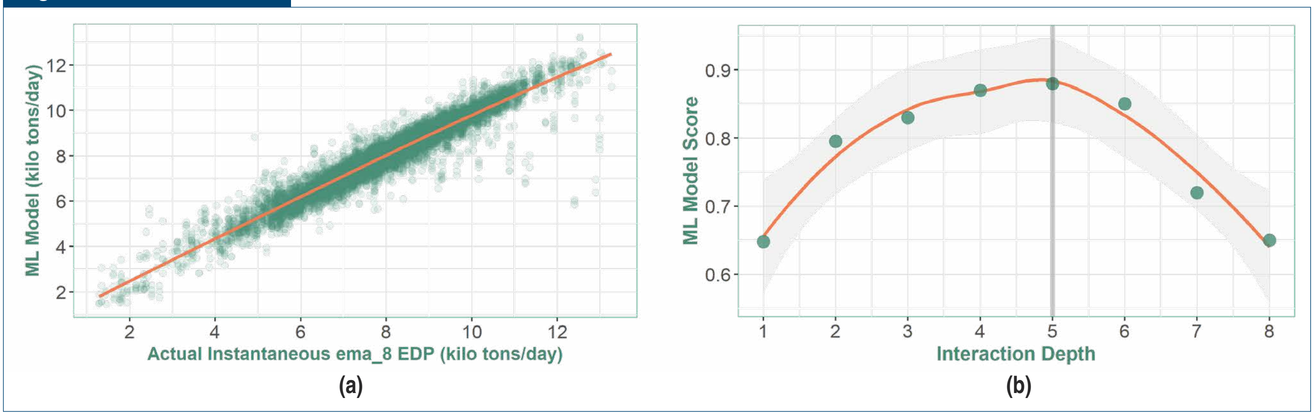

ML model vs. actual instantaneous ema_8 EDP on test subset for best model score (a) and interaction depth = 5 provides the best model score (b).

there are only 13 terms that are necessary to explain the variance in the data; these terms carry eight features (called variables of importance or varimps) with four-way interactions. The top five of these variables of importance which were listed in Fig. 2 showed that K is the most important variable and literature supports this find. That K ranks at the top of the list is not altogether surprising and is a testament to its role as a pathway or a proxy variable; the ML model now allows one to quantify its impact more precisely.

Interaction Depth — The blast furnace is a complex process lending itself to deep interactions, which simply means several variables act together to influence K and as a result influence hot metal output, so that when some conditions in several variables are met simultaneously, the hot metal output is impacted more severely. The overall depth of these interactions can in fact be measured and it is shown in Fig. 7. In fact, one of the reasons why traditional methods of statistical analysis and model representation fail is because of significant deep interactions that remain unaccounted within a vast space of possible interactions. Fig. 7 suggests that at least three-way cross-interactions are necessary to include in the model for sufficient predictiveness and up to five-way interactions. Including higher than five-way interactions offers no advantage; and adding more than six-way interactions begins to catch too much noise in training to degrade the model severely when measured against the test subset. It has been concluded that the blast furnace process is represented ideally by a five-deep interaction and sufficiently by a three-deep interaction.

With 115 predictors and a five-level deep interaction, this gives more than 150 million possibilities. Of course, only a handful of terms and interactions are meaningful — but at the outset, in an unbiased model, it is not known which of these are relevant. If a standard non-linear regression model is used that includes cross-terms, with these many coefficients, an equivalent number of data points is required and typically 10 times larger for an adequate confidence. This would require 1.5 billion points. Clearly, that much data is not available, or technically speaking, the data set is not powered. Additionally, curse of dimensionality limits discovery in high-dimensional sparse spaces. All of this means that an optimized way of extracting variables of importance that influence EDP is necessary and all of this from 4,400 data points in the data set. ML is that optimized method that overcomes the limitations emerging from a vast "naïve" space that includes interactions.

Top Variable — Fig. 8 shows permeability and EDP as a function of lime percentage in sinter. The result is somewhat counterintuitive; it is a reasonable expectation that a higher percentage of lime will result in lower production because lime in sinter reduces the iron content. However, it turns out that K continues to drop (i.e., permeability continues to increase) and EDP continues to increase as lime percentage increases from 10.5% peaking at 11.8%. It is believed this is due to the binding action of lime in sinter,[2] allowing it to stay intact in the blast furnace by increasing the strength of sinter. Sinter with higher strength is preferable because this requires less lime percentage to enhance binding and inhibit mechanical degradation; but for sinter of a given quality, there is an optimum level of lime percentage to maximize permeability and hot metal output.

Figure 8

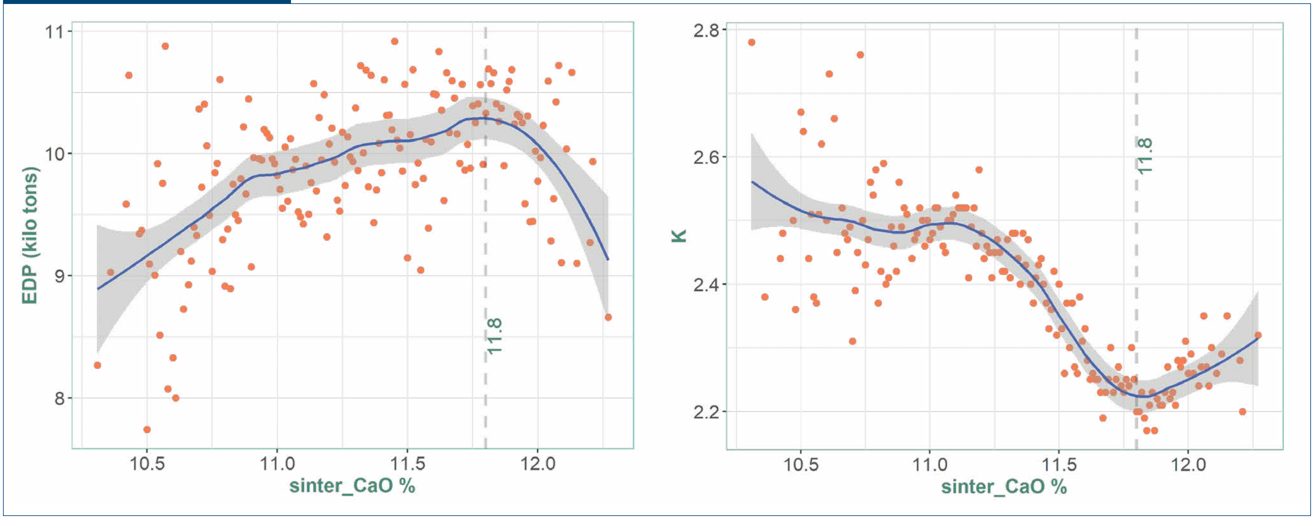

Sinter lime effect on permeability and production. (Lower K = higher permeability.)

Confounding/Causal Inference/Association — An important note needs to be made regarding the methodology. The data in this analysis is observational and not acquired through a designed experiment. It is hard to do a designed experiment in blast furnace because of the importance in keeping the plant running (downtimes are generally not acceptable in the normal course of business). However, drawing conclusions from observational data poses some constraints; in particular, confounding can be present in the data. For example, as the lime percentage in sinter changed in the data set at various time segments in the data set, one should ask whether some other parameter also changed to such a degree at the same time (for completely unrelated/independent reasons) such that potentially this other parameter could have caused the effect noted earlier on permeability and hence on EDP. Alternatively, the following questions are posed: Is sinter CaO% causally related to EDP and K? Or is it merely a non-casual association in the data set?

To answer this, Bayesian thinking and counterfactual analysis are used. The authors examined the biggest factor that could potentially be confounded temporally with sinter CaO% and be the source of the effect. Coke quality is known to be a significant contributor to permeability through the M40 measurement, which is a measure of coke strength. Higher M40 implies better mechanical strength of coke, which reduces fines in the column; fines act to block the free movement of gas by plugging interstitial spaces. To examine the effect noted in Fig. 8, and possible confounding, the next section compares a time period when permeability and production were the highest against a time period when they were the lowest during the evaluation period.

Good vs. Bad Production Periods — In addition to examining the confounding noted earlier, good vs. bad (period) analysis — also called GVB analysis, serves another purpose. Metallurgical and other specialty engineers know the domain extremely well but may have only a passing familiarity with data science; this demographic can derive immense benefit from GVB visual charts, making it easier to disseminate than

Figure 9

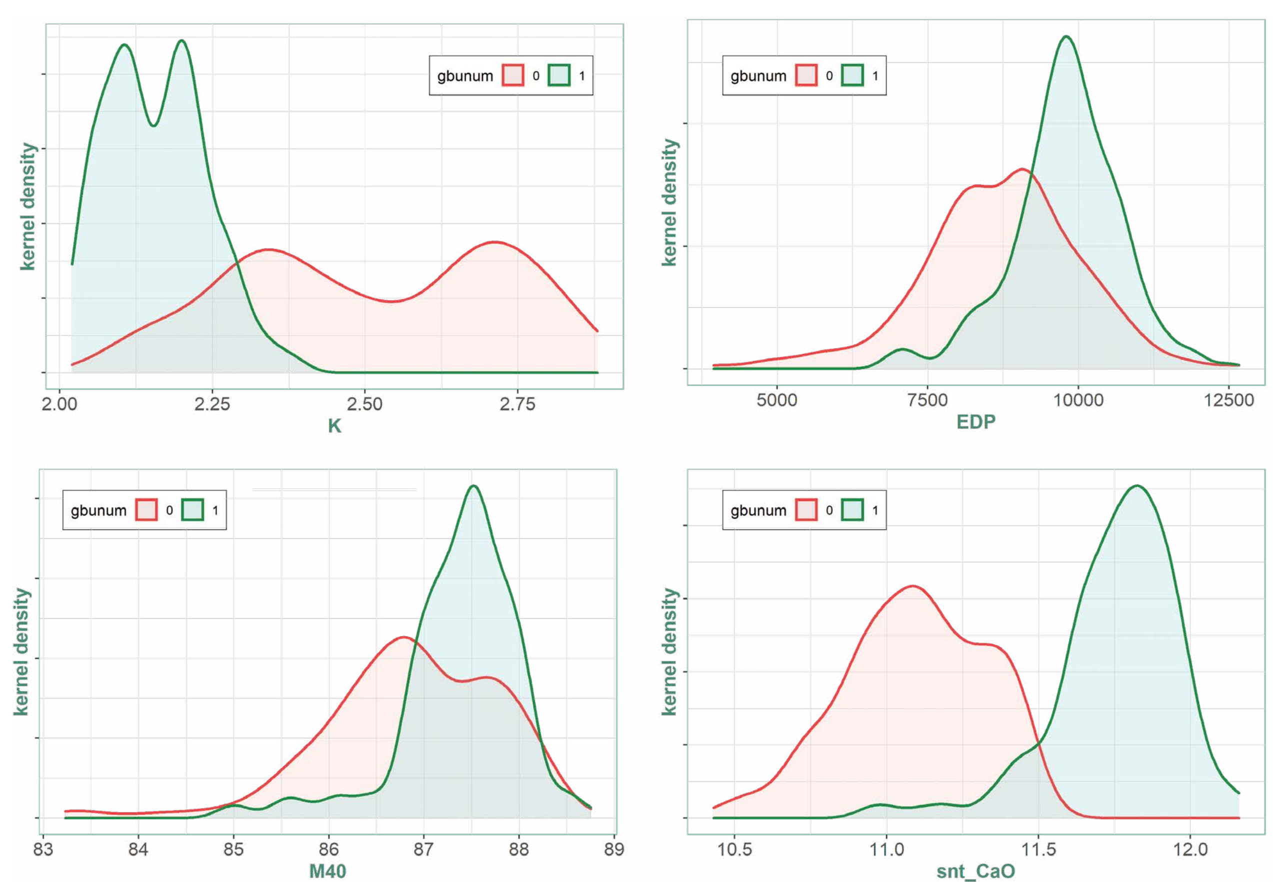

Good vs. bad (GVB) univariate charts for EDP, K and sinter CaO% and coke M40 for good production high-permeability period (green) and bad production low-permeability period (red). While M40 explains part of this difference, sinter CaO% explains more.

more complex and esoteric ML interpretations. So, the key varimps are shown in two extreme conditions in Fig. 9 — a period where the permeability was at its highest, i.e., K was its lowest (in green) and a period where the opposite was true (in red). Of course, as stated earlier, the ML model (Fig. 7) includes complex interactions that are not visible in Fig. 9, but univariate charts (one variable at a time) are easy to disseminate and therefore still retain significance in practical applications even when complex ML methods are used to develop the production model.

Fig. 9 shows that while M40 explains part of the difference in good vs. bad periods, as expected, an equally influencing factor in the model is sinter CaO%, as visible in the starkly separate distributions. To examine counterfactually, the effect of M40 is controlled by restricting it to less than 86.5, in Fig. 10, effectively examining the subset of data where coke strength was bad in both good and bad periods of permeability. With this additional condition controlling for the effect of M40, the sinter CaO% difference in distributions is even more stark! M40 and sinter CaO% independently influence permeability such that if permeability is impaired due to incoming quality problems in coke, then a higher effective sinter strength can partially compensate for it.

Effects of other varimps emerging from the model, such as coke moisture and ash and several others, are as expected based on prior domain knowledge; therefore, discussion of it is not included here. Instead, the focus is on those areas that add value to the existing body of knowledge. An astute reader will want to pose the question whether a model with this much accuracy as is presented here could have been determined without invoking ML. The short answer is no; among other reasons, not enough data to fit 150 million coefficients, let alone determine them with a high enough confidence. Model accuracy after n-fold cross-validation and regularization shown in Fig. 7 was only achievable through ML and incorporating interactions in the model at the right level of depth.

Optimization Criteria — As mentioned earlier, ML has distinct advantages over traditional methods of data modeling when there are deep interactions among many features. As a result, ML models can address vast feature spaces. However, this comes at a cost — ML models must be updated every so often. Natural phenomena do not change over time, of course; however, the number of possible combinations of inputs that impact the output is large and only some of these combinations are impacting the process at any given time and are therefore included in the model at the time it is built.

In general, it is preferred to have minimum complexity models that are above an acceptable ML model score. However, increasing the number of terms beyond the maximum ML score model could potentially increase the lifetime of the ML model. This is somewhat counterintuitive because while increasing model width includes more noise, it also allows inclusion of terms that are impaired compared to the strongest terms currently but could well become prominent with slight shifts in inputs. In a future data set, it is likely that some of these additional terms could be the strongest. This wider model that includes more terms (and more noise) is preferred by some practitioners to extend model lifetime at the expense of a somewhat diminished model score while still above a minimum level. Most practitioners, though, prefer the highest

Figure 10

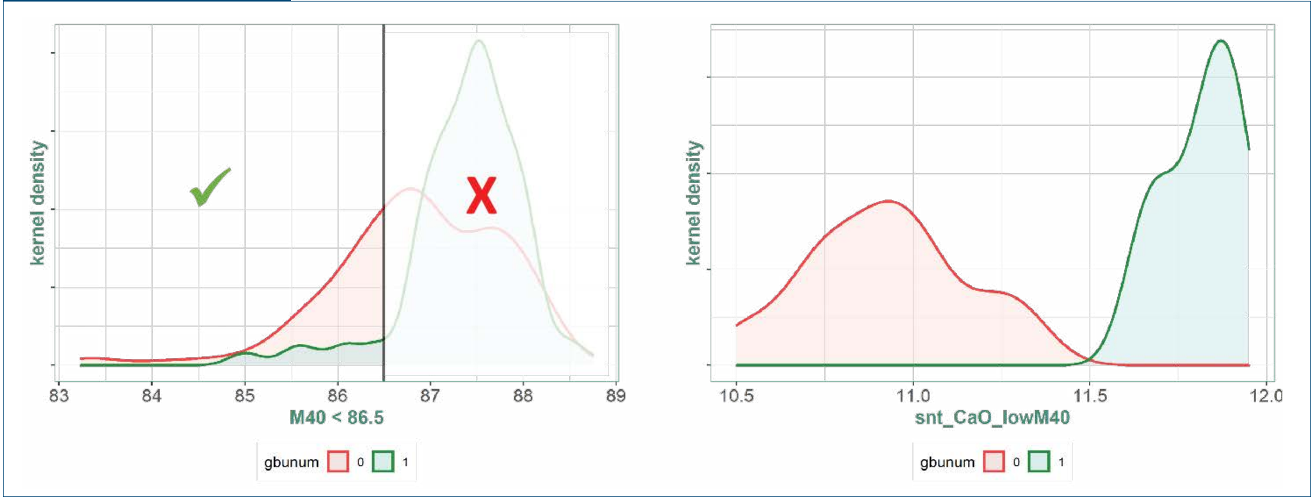

Controlling for the effect of M40 on K and EDP, by restricted to <86.5 in both good period (in green) and bad period (in red). The distributions of snt_CaO% for good and bad periods are now even more starkly separate.

Figure 11

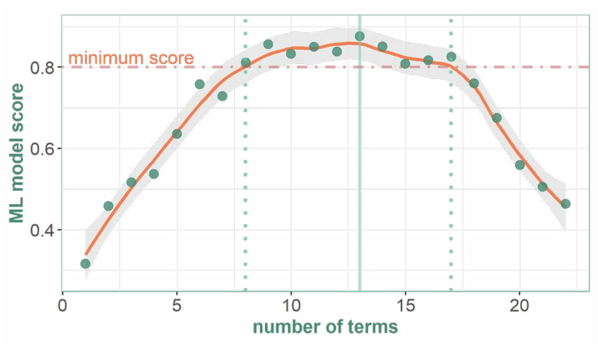

Selection of number of terms based on optimization criteria.

ML model score operating point since model updates are getting inexpensive to compute nowadays with automation.

Note on Modeling Culture(s) — Lastly, the authors leave the reader with a philosophical pondering that Leo Breiman raised in the much-quoted paper from 2001.[5] He argued that there are two cultures in what is now called the analytics domain — (a) data modeling culture, for which the term physics-based models in this paper has been used and (b) algorithmic modeling culture which is the province of ML and deep learning, etc., but divorced from a physical understanding of the process. The authors do not offer a leaning in either direction but only suggest that awareness of this user bias, and by implication, blind spot(s), is helpful in selection of methodologies and even interpretation of results during an advanced analytics exercise, whether addressing an out-of-control condition or continuous improvement or real-time prediction of production catastrophic events.

Summary and Conclusions

This paper has shown that productivity improvements in a blast furnace can be meaningfully assisted with ML methods with cross-validated and regularized accurate models displaying good reliability. Key variables of importance were ranked, scored and explained based on physics as well as with algorithmic interpretations. It was shown that permeability metric (K) acts as a proxy variable for production; in turn other variables of importance drive permeability, including coke strength, percentage of lime in sinter, coke ash and coke moisture, among others. And the physical and algorithmic rationale for the counterintuitive finding of CaO% in sinter as one of the variables of importance was explained.

ML can be used not only for blast furnace productivity and performance modeling and improvement but also by any other department in the steel plant. The authors posit that ML can account for the vast degree of variability found in blast furnaces around the world, thus providing guidance to those that operate below benchmarks.

The next evolutionary steps in applications of ML and deep learning to iron and steel manufacturing should include near-real-time applications such as early warning signal for blast furnace slip or other catastrophic unscheduled events that are immensely expensive for the steel plant. While the primary technology that was exploited in this paper falls under the umbrella of supervised learning for continuous improvement and process quantification for meeting strategic long-term goals, newer methods in unsupervised learning as well as deep learning can be of great tactical assistance for production monitoring and real-time "imminent failure" detection.

Acknowledgments

The authors would like to thank the entire staff at the blast furnace operations at the JSPL plant at Angul, Odisha, India, and in particular the senior management of iron zone and blast furnace operations: Damodar Mittal and Arvind Bhagat for making resources available to make the work possible. Technical and enlightening discussions with Paritosh Haldar were invaluable. Executive guidance from Sudhanshu Saraf and Pankaj Gautam were critical in project progression and the authors thank them. Promotion of this project by Venkatesh Jindal was immensely useful for getting the project initiated and the authors are grateful for that. Finally, the authors thank Charanjiv Singh, principal at Enabling Excellence, for setting up this project.

References

1. S. Natsui, S. Ueda, H. Nogami, J. Kano, R. Inoue and T. Ariyama, "Analysis on Non-Uniform Gas Flow in Blast Furnace Based on DEM-CFD Combined Model," *Steel Research International,* Vol. 82, 2011, No. 8, p. 964.
2. R.P. Bhagat, *Fundamentals of Iron Ore Sintering,* Chapter 11, p. 146.
3. B. Boehmke, Greenwell, Brandon, 2019, *Hands-On Machine Learning With R,* Chapter "Gradient Boosting," Chapman & Hall, pp. 221–245.
4. T. Chen and C. Guestrin, XGBoost: A Scalable Tree Boosting System, https://arxiv.org/pdf/1603.02754.pdf.
5. L. Breiman, "Statistical Modeling: The Two Cultures," *Statistical Science,* Vol. 16, 2001, No. 3, p. 199. ✦

This paper was presented at AISTech 2021 — The Iron & Steel Technology Conference and Exposition, Nashville, Tenn., USA, and published in the AISTech 2021 Conference Proceedings.