Adopting Machine Learning Based Workflows for Reducing Production Risk and Cost

Pamir Ozbay[1], Carter Almquist[2],

[1]Fero Labs
433 Broadway Street, New York, NY, USA, 10013
Email: pamir@ferolabs.com

[2]Gerdau Petersburg
25801 Hofheimer Way, Petersburg, VA 23803
Email: carter.almquist@gerdau.com

Keywords: machine learning, melt shop improvement, chemistry refinement, optimization, software implementation

INTRODUCTION

The integration of machine learning (ML) into a steel mill production workflow is one of the most challenging aspects of ML-driven process improvement initiatives. Most discussions on ML center on early-stage issues such as data availability or quality. These are important considerations as the performance of ML models is dependent on training data. However, overcoming these modeling challenges is only the first step. The next, and more challenging phase, is the implementation.

Since 2020, this steel mill has been using ML software to optimize the chemistry of grades at the melt shop in real-time. Initially, the potential of using ML was demonstrated through offline tests and simulations. Full benefits were achieved by connecting live production data to software and with active collaboration between experts across departments. On any given day, ladle furnace operators use the ML software's recommendations to adjust alloy composition for each heat. The improvement facilitator maintains operator training and ensures key learnings are disseminated across teams. The quality group maintains the ML models and updates constraints given to the models to guarantee quality targets are met. The advanced automation features and intuitive interface of the software have been critical in facilitating coordination, streamlining training, and reducing model maintenance workload. These combined efforts have driven widespread adoption of ML in steel production, resulting in a substantial increase in profitability.

Since live implementation of the ML software, thousands of heats have been cast with ML optimized chemistry. This has led to an average reduction of $3 per ton in alloy costs, while maintaining stringent product quality standards. The software has also been successfully scaled to the other mills of the producer. Furthermore, the production of ML optimized steel has not only reduced alloy costs but also had a positive impact on the environment. The decrease in unnecessary alloy additions has resulted in the conservation of over 500,000 pounds of manganese, carbon, niobium, and vanadium, reducing the demand for their mining and transportation.

In this paper, we present a comprehensive overview of the strategies and tools used to address process and quality challenges in a steel production setting. We first introduce the traditional practices, then delve into the critical software and technology features that have been instrumental in the successful implementation of ML models. Finally, we offer a set of practical recommendations based on our insights and experiences, aimed to aid other organizations in their own implementation efforts, followed by results achieved at this mill.

MAXIMIZING PRODUCTION QUALITY AND EFFICIENCY

Operational challenges at the melt shop

The resulting chemistry in a heat of steel is heavily dependent on the stability of the current operating conditions. The content of many residual elements such as nickel, copper, and molybdenum depend on scrap mix and scrap quality charged into the electric arc furnace (EAF). Furthermore, plant operations are fluid and the alloy content from heat to heat can change based on operator actions and the reactivity of the heat. Given the variation in chemistry and the downstream rolling conditions, planning for the final mechanical properties of rolled steel at such an early stage in the process is a complex challenge.

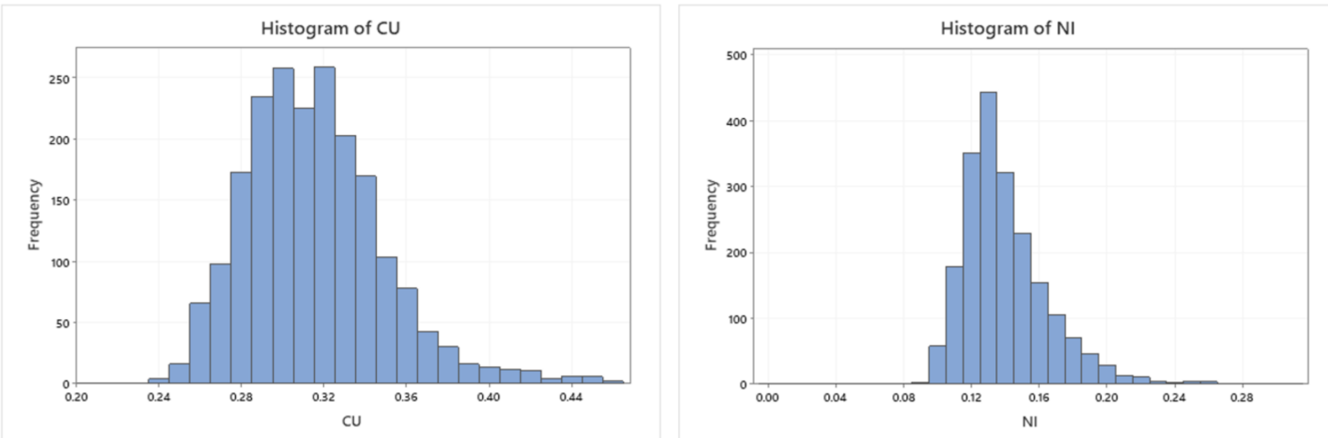

Figure 1: The residuals from each heat exhibit considerable fluctuations, which exacerbates the impact on the quality of the final product given the variability in the overall production process. To achieve quality goals, the quality control team must meticulously outline operational limitations and regularly revise procedures to keep the constraints aligned with recent changes in production.

In addition to chemistry that is in constant flux, the process conditions in the melt shop also introduce variability with direct impact on final product quality and operational costs. These conditions include:

- Scrap quality: not only impacts residual elements, but also can affect EAF target deliverables such as tap oxygen content, slag chemistry, and processing time which have significant downstream implications.
- Coordination: good communication between the EAF and ladle furnace (LMF) operations can help minimize chemistry composition changes from heat to heat.
- Slag content: suboptimal slag impacts alloy recovery, desulphurization, and reaction rates at the LMF. Slag content can be heavily impacted by the amount of carryover slag from the EAF (and EAF slag content).
- Alloy recovery: unplanned or delayed alloy recovery due to aforementioned sources of process variability can misinform operators and misguide their decisions.
- Alloy costs: the costs of the alloys used in the melt shop varies throughout the year based on market conditions and procurement practices.
- Spectrometer noise: measurement and calibration noise of the spectrometers used to determine steel chemistry adds an additional source of error and therefore reduces information quality.
- Production goals: timely action is a must at the furnace as the operators need to account for all the above conditions while ensuring a constant flow of material to the continuous caster.

Existing methods of data driven approach to account for the challenges at the melt shop

A key approach for accounting for such variation is capability analysis. This is a statistical method used to evaluate the performance of a process relative to key specification limits. The process involves collecting data on key process variables, such as dimensions, chemistry, or mechanical properties and analyzing the data to determine the mill's capability[1]. A typical workflow is as follows:

1. Define the process: The first step in conducting capability analysis is to define the process that is being evaluated, including the inputs, outputs, and key process variables.

2. Collect data: The next step is to collect data on the key process variables. This data is typically collected over a period, such as a day or week, and should include a representative sample of the products produced by the process.
3. Calculate process statistics: Once the data has been collected, process statistics such as mean, standard deviation, and range are calculated for each variable. These statistics are used to determine the process's capability.
4. Determine process capability: The process capability is determined by comparing the process statistics to the specification limits for the product. Capability indices such as Cp, Cpk, Pp, and Ppk are calculated to measure how well the process can meet the specifications.
5. Analyze results: The final step is to analyze the results of the capability analysis and identify any areas where the process is not meeting specifications. The aim is to identify the root cause of the problem and implement corrective actions to improve the process performance.
6. Continuously monitor: It is important to continuously monitor the process performance and re-evaluate the capability of the process after any changes have been made. This will help to ensure that the process is always operating within the desired specifications.

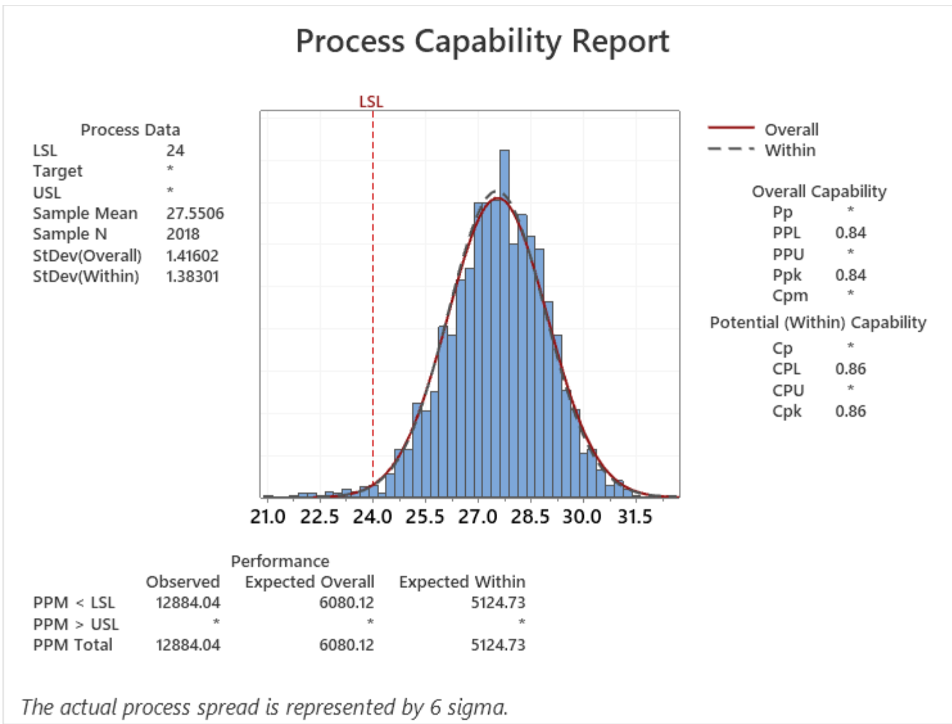

Figure 2: An example process capability analysis performed on a subset of the production (data not from the mill).

As the steps above highlight. Capability analysis is a valuable tool for evaluating the performance of a process at steel mills, but it does have some limitations:

- Assumes data collected follows a normal distribution, which may not always be the case. This can lead to incorrect conclusions about the process's capability if the data is not normally distributed.
- Limited to measuring existing performance rather than identifying areas for improvement and predicting future performance.
- Limited to a specific set of process parameters and may not take into account other important factors that can affect the overall performance of the process.
- Does not account for future variability of raw materials or process variation.
- Requires manual data exports to be collected, which can be time-consuming and expensive.
- May not account for all process variations in the process, especially when dealing with complex and multivariable processes, leading to incomplete conclusions.

Machine learning based data driven approach to account for the challenges at the melt shop

The main objective of machine learning (ML) is to use past data to uncover patterns and relationships. Unlike traditional modeling methods, ML does not require the user to explicitly define the relationships (such as the function linking tensile strength to mill parameters). Instead, ML uses historical data to "learn" and determine the function on its own, making it a data-driven approach to modeling relationships.

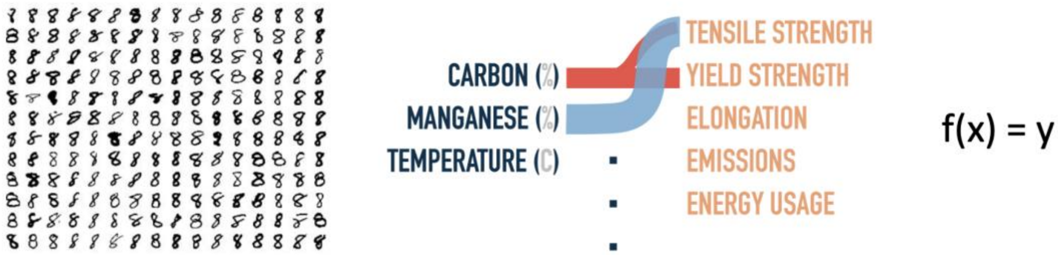

Figure 3: ML models use production data learn the non-linear functions that map process metrics to the target KPIs. These functions are continuously updated with new production (i.e., new data).

ML has distinct advantages for addressing operational challenges in steel mills. Traditional methods would struggle to analyze large amounts of historical data with hundreds of tags and multiple years of heat data. Even if this were possible, traditional tools are not equipped to stay current with the latest changes. It would be impractical for a mill to perform a capability analysis every week, but the plant's conditions can change significantly even within a day. The ML-based data-driven approach provides flexibility to adapt to these changes. Furthermore, the non-linear models map the complex relationships more accurately, leading the better insights and recommendation.

Focus use case for this paper:

Minimize alloy costs in real-time at the ladle furnace through early prediction of final mechanical properties

With the white-box ML model, the mill can accurately predict key final quality tests, such as yield and tensile strength, in real-time. The model is trained on a comprehensive data set that includes all of the mill's products and incorporates data from the melt shop, rolling mill, and test lab. This data is updated in real-time through a live connection, allowing the model to be automatically retrained within minutes, incorporating the latest information and insights in its recommendations.

The ML model is used to render early predictions of the final rolled product while the steel is still in the melt shop. The software takes these predictions and the mill's quality specifications into account, providing alloy addition recommendations after each ladle sample. The operators follow these recommendations to save on average $3 per ton of steel. The software also generates additional predictions after the heat has been cast, providing the mill with a final check of the forecasted mechanical properties before the billets are released for rolling.

While alloy reduction is the most popular use of ML within this producer’s mills, the other use cases the mill’s process and quality experts have explored include defect prediction, tap oxygen forecasting, reheat furnace emission reduction and energy optimization. The common goal across these use cases is to leverage ML technology to continue to streamline different processes, improve quality, and optimize costs.

Given the novelty, considerable time is spent on ML algorithms and different modeling techniques. This paper will not delve into the particulars of data science methods. The goal is to highlight some of the key features that have been found beneficial by the mill operations to thoroughly implement software that delivers new capabilities resulting from the use of ML.

Important ingredients that increase use of ML by mill operations

Many potent technologies suffer from low end-user adoption. Complex and novel technologies such as machine learning have proven to be even more vulnerable on this front[2]. To avoid this, change management challenges need to be tackled from the start alongside data and model related activities. This helps stakeholders understand and adapt to the new technology and its impact on the organization and ensures smooth integration of machine learning into existing processes and systems. Furthermore, it minimizes resistance and disruption to the organization following an initial pilot period and ultimately, increases the chances of success and return on investment of the implementation.

The experience gained improving adoption of the ML software by the mill operations has highlighted the importance of:

- User-friendly interface and easy integration
- Easy-to-maintain models
- White-box machine learning

It is worth emphasizing that these features do not negate the need of having good quality data and high-performance models. However, without these features, even the best performing ML models risk being never truly used by the mill operations.

User-friendly interface and easy integration

Having a user-friendly interface is critical. The ML models may be accurate and offer important insights, but this will not translate to workflow of profit improvements if employees do not use the new tool. Furthermore, overall product experience has very significant impact on user productivity and ultimately, the mill's profitability. During initial phases of the live deployment three patterns emerged as frequent challenges. Addressing them were crucial in driving adoption.

1) There is limited time available for training and for a complete implementation. All shifts need to be trained in the proper use of the new technology. Reducing the learning curve allows improvement facilitators to easily train operators.
2) Machine learning is complex. The software needs to distill the output of ML models as visual and interactive dashboards that anyone at the mill can understand and use. Streamlined dashboards also have the added benefit of reducing the likelihood of user errors.
3) User-friendly interface improves employee satisfaction and engagement with the software, which leads to greater adoption and use of the software in the long term.

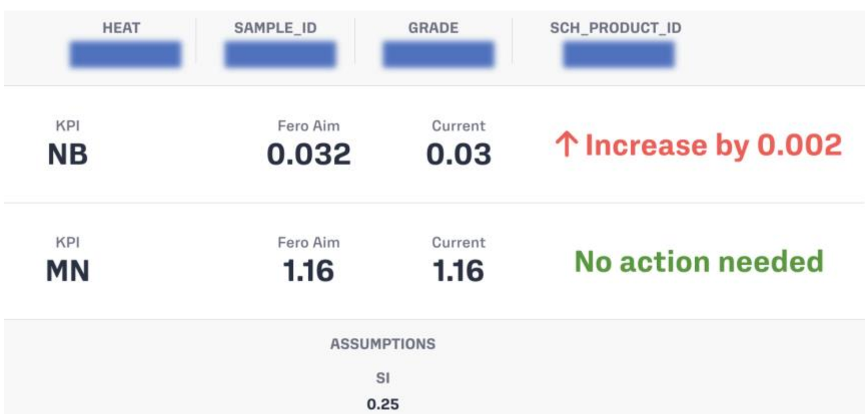

Figure 4: Software must translate the ML model to actionable recommendations for quick decision making. This figure is the ladle furnace operator dashboard of a live white-box ML model. The model is trained with over 20,000 heats representing the entire grade portfolio of the mill with tags from the melt shop, rolling mill, and the test lab.

Even a streamlined interface may prove to be difficult to adopt. This can be due to limited real estate in the control rooms or a shortage of training resources. To overcome this, ideally any new software added to the technology stack of the operations should be able to connect to the readily available software. In the context of this use case, this is accomplished through a standard REST API interface. The mill's IT team uses these end points to transfer the outputs of the ML models directly to the HMIs in the pulpits. While the Petersburg operations team preferred to use the UI of the software, there are other mills that have adopted a hybrid approach. The key is to find a solution that fits the specific needs and preferences of the operations team, while still providing the critical insights generated by the ML models.

Easy-to-maintain models

Maintaining ML models is important for several reasons. One reason is to ensure that the model continues to perform well on new data. As the distribution of the data can change over time, the model's performance may degrade if it is not retrained or updated. Additionally, maintaining the model can help to prevent overfitting, which occurs when a model becomes too specialized to the training data and performs poorly on new data. Finally, maintaining the model can help to ensure that it remains fair and unbiased. This includes addressing any issues of bias in the training data and monitoring the model's performance on different subsets of the data to ensure that it is not making unfair predictions[3]. These points become especially valid when a natural drift is present across a large set of process and raw material measurements.

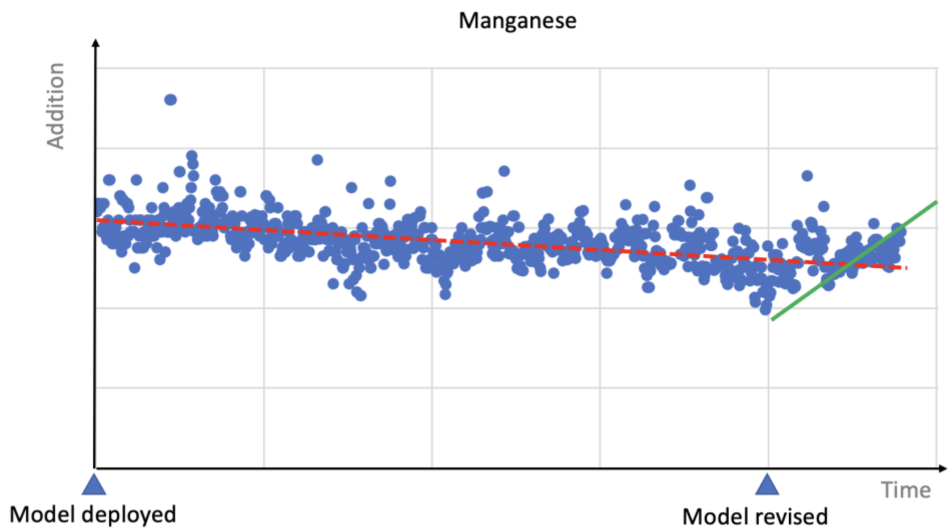

Figure 5: The frequency of model maintenance is important as well. A model that is not maintained would have missed the long-term downward trend (red) following deployment. If it also not retrained regularly, it would also miss the more recent upward trend (green).

On average, maintaining high model performance reduces value erosion by 60% over the course of a year[4]. Since retraining and validating a model can take up to a day of qualified staff's time, this activity almost never receives the regular attention it deserves.

Approach	Maintained by	Data Scope	Frequency	Time Needed
Process capability – Six Sigma	Process experts	Selected subset	A few times a year for a limited subset of the production.	2 hours per product—can only focus on one product at a time to minimize process variables.
Bespoke ML model	Data scientists	Entire production	Highly dependent on the availability of a very limited pool of qualified resources that at minimum includes a data scientist and a process expert.	6-20 hours per update, dependent on the expert availability and issues introduced since last update
Automated ML software	Process experts	Entire production	If desired, can be done multiple times per hour thanks to automation. The typical cadence is once a month.	Less than 10 minutes per week based on usage patterns at the Petersburg mill.

Figure 6: Comparison of effort needed to maintain different statistical approaches available for process improvement.

Beyond keeping performance above the threshold, the live deployed ML models at the mill have been improving their performance since 2020. The optimal retraining frequency during typical production has been to be around once a month. This frequency can be increased to retraining every few days especially if the mill has introduced new grades or is going through a unique period such as ramp-up following a shutdown. Since retraining is automated by the software and takes minutes vs. days, metallurgists and engineers can perform this activity without requiring any additional outside help. Furthermore, this allows these experts to keep a close eye on the trends and note any unusual changes following the retraining of the model if these models are white-box.

White-box machine learning

Even with the most accurate ML models, the mill's operation team would not have adopted the new technology if it used black-box models. White-box machine learning is critical in industry as it allows for transparency and understanding in the decision-making process of automated systems. The users at the plant use these insights to apply their domain expertise to address any potential biases or errors in the model. This creates a virtuous collaboration cycle that improves the models, increases the trust of the users in the new tool, and ultimately the use of the new technology in daily routine.

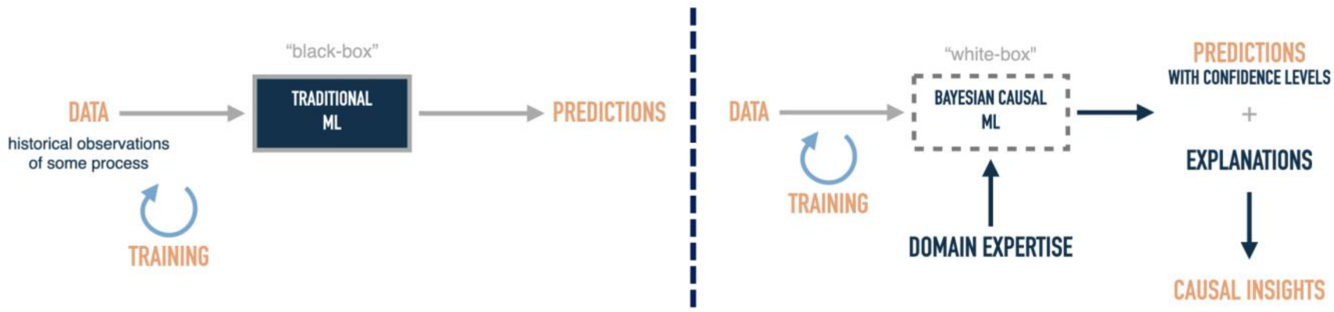

Figure 7: Our experience has repeatedly shown that without white-box techniques, the process experts do not trust the recommendations of the ML models. A new tool will never get scaled or used regularly if the users do not trust it.

Another critical feature of the white-box model[5] is the confidence bands. This is how the software conveys the statistical confidence of its predictions and recommendations to offer an additional tier of transparency. There are production zones where the ML models are not as accurate, and it is important for the software to make the operations aware of the prediction made in these zones. The confidence bands in the deployment tend to expend when one or more of the following conditions are met:

1. When introducing a new grade or producing a product that is seldomly rolled, the models may enter a low data zone. Data driven approaches such as ML do not excel in such zones.
2. The historical training data may have quality problems due to noisy sensor measurements or data logging practices. The models need to be robust to handle this noise and make the user aware of such issues by conveying low confidence in its predictions when necessary.
3. There are products that are more prone to phase transformation due to unique physical characteristics. In such cases, the similar chemistry and rolling practices may still result in very distinct mechanical properties. This will expand the confidence bands as there is a variation in the production process that is difficult to control.

Ultimately, some grades or days of production may prove to be too risky or difficult to optimize with ML. This is exactly the behavior observed at mill. When the residuals coming from scrap are unusually low, the ML software may struggle to confidently recommend that production will meet mechanical specification requirements. The software also automatically rules itself out in difficult to produce grades where the source of the difficulty goes beyond the chemistry. This does not impact the operations thanks to smooth integration of the software; the operational constraints set by the quality team is always readily available as a fallback measure when needed.

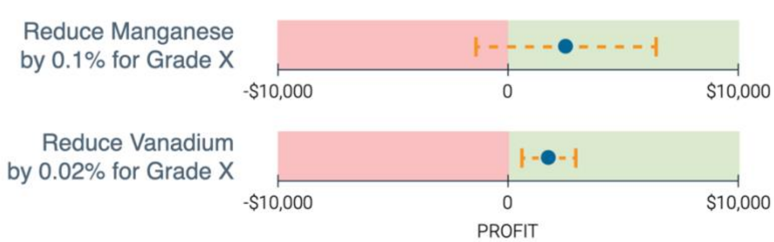

Figure 8: Confidence bands have massive implications on how the ML model's recommendations are used in real world. This figure represents two recommendations with roughly the same mean prediction point (blue dot), and yet with significant difference in confidence band (orange line). In this instance, while the Vanadium recommendation meets the mill's requirements, the manganese suggestion is ignored. The constraints of the requirements can be adjusted on the fly for maximum flexibility and minimum risk.

KEY LEARNINGS

Operational best-practices for maximizing the benefits of using machine learning

The learnings highlighted in this section will ensure that ML tools are effectively integrated into the operational workflow, and recommendations from the models are accurately monitored and acted upon. By following these best practices, the mill's operations can reap the full potential of an ML tool faster, drive substantial improvements, and scale the benefits.

Focus on training and collaboration

The effectiveness of ML software is dependent on the proficiency of its users. To maximize the benefits of the new tool, training of its daily users such as operators and engineers was prioritized. The features highlighted in the previous section facilitated a smooth transition to the new workflow by reducing the learning curve and building trust in the software's recommendations. During training, a close collaboration between the quality department and the melt shop improvement team played a crucial role in successfully onboarding new users and expanding the application of the new tool.

Our suggestion is to identify key stakeholders from the operations and quality teams early on and to engage them throughout the deployment process. Implementing high-accuracy ML models without involving these users will not result in widespread adoption within the mill. To ensure the successful and frequent use of the new technology, it is especially important to involve the operations team at every step of the process. It is crucial to address the following questions at the outset of the project to involve the appropriate stakeholders and establish an effective user adoption framework before the software deployment phase:

1. What is the current problem or challenge we are trying to solve?
2. Who will be impacted by this change and how?
3. What are the potential benefits of this change?
4. What are the potential risks and drawbacks of this change?
5. How will we measure the success of this change?
6. What resources (time, money, people) will be required to implement this change?
7. What is the timeline for implementing this change?
8. How will we communicate and engage stakeholders throughout the change process?
9. What support structures will be put in place to help people adapt to the change?
10. What is the contingency plan in case the change does not go as planned?

Make the software physically accessible

The simplest modifications can often have a significant impact. A great example of this was the activation of the operator interface at the ladle furnace pulpit and running it continuously on a centrally positioned monitor. This change led to a 2.5-fold increase in the tonnage of steel cast using ML recommendations and a 3-fold improvement in profits in just one month.

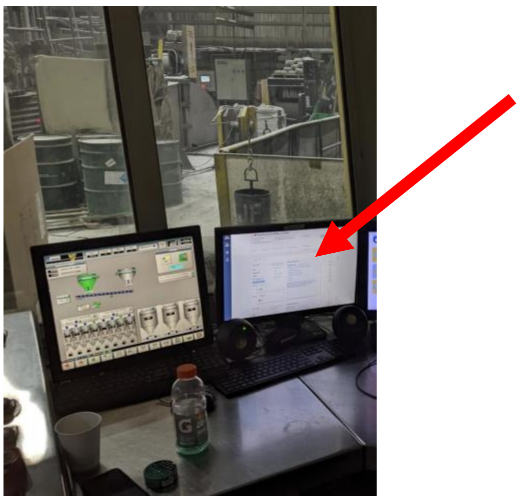

Figure 9: Designating a centrally positioned monitor in the control room is an effective way to enhance the utilization of ML in production.

Prioritize production volume over aggressive changes

It is better to use ML with as much of your production as possible even if this requires being more conservative than necessary. We accomplished this by adding artificial buffers to the actual specification limits. This reduces the dollars per ton savings potential, in some cases by up to 55%, but the loss is easily made up with the increased volume of production using ML optimized recommendations. The quality team of mill has the ultimate control on the buffers. These buffers are adjusted periodically at the grade level to minimize risk while increasing profit.

Using buffers also has the added benefit of reducing the ramp-up time the operations team needs. If the use of the tool is contained to limited trials, the users do not build up the operational muscle memory needed to best leverage the new workflow. Once the software has been used conservatively for 100s of heats and the operators have gained experience, the buffers can always be relaxed to improve dollars per ton savings.

The results detailed in the last section of this paper show that, despite the conservative approach described in this section, the alloy addition has been reduced by approximately 8%. The true potential savings is twice this amount, but the mill wisely adopted a gradual approach. This gives users the chance to adjust to the new workflow. Equally important, an incremental change also provides additional data to ML algorithms in the low alloy operating zone that further enhances the prediction accuracy and minimizes the need for extrapolation.

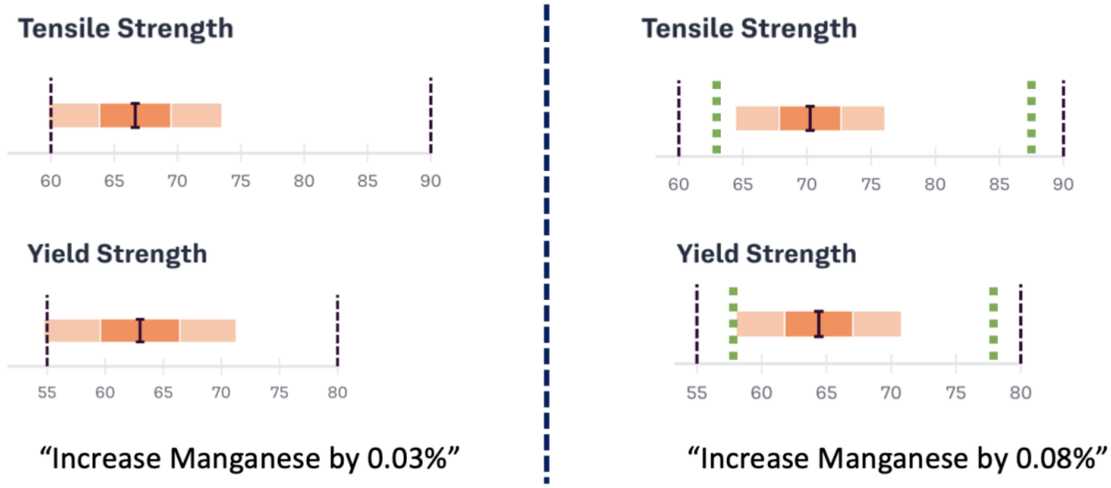

Figure 10: White-box ML helps with reducing risk by utilizing confidence bands. In the example on the left, the software is configured to match specification limits (black dashed lines) with at least 95% confidence. The software can be forced to be more conservative by adding a buffer (green dashed line) to the official limits. The example on the right showcases how the recommendation changes for the same heat when the buffers are introduced to guarantee safe production.

Increase operational flexibility

The implementation of ML software to optimize raw materials and process parameters results in a substantial reduction in the variation of the final mechanical property test results, with an improvement of up to 66%[6]. This enhances the stability and quality of the final product. However, achieving this does require increased operational flexibility during production.

In the absence of ML models, the quality team employs capability analyses to calibrate the grades to ensure the mill meets the product quality targets. To guarantee a low failure rate, these grades take into account the worst-case production scenarios that may occur due to below average alloy content and suboptimal rolling mill conditions. It is expected that, under normal operations, the rolled product will meet the required mechanical specifications as long as the operators stay within the grade specific alloy addition range. The recommended alloy addition range in the grade book is stable and updated only every few years. For example, if the mill is producing a single grade for a full day, the operators do not need to continuously consult the HMI to know the recommended alloy range after the initial few heats.

With ML recommendations, the operators receive tailored recommendations based on the heat's chemistry and the process metrics associated with the rolled product. The aim of these recommendations is to minimize alloy addition costs, as long as the software predicts that the quality targets will be met with high confidence. This approach differs from designing a grade for the worst-case scenario. In real-world deployments, the software may recommend reducing alloy additions when it detects that other strengthening elements are sufficient to meet the target mechanical properties or increase alloy additions when it predicts high risk to final quality. As a result, the dynamic targets fluctuate from heat to heat, requiring the ladle operators to pay close attention to the latest alloy content recommendations and optimize the heat's chemistry. To maximize cost savings, the main furnace operators must also adjust their workflow, ensuring the initial tap additions leave room for the ladle operators to add up to the recommended values. For instance, if a heat is shipped to the LMF with a 1.20% Mn, but the software would only recommend that 1.17% Mn is needed to remain within the constraints, this heat cannot be truly optimized.

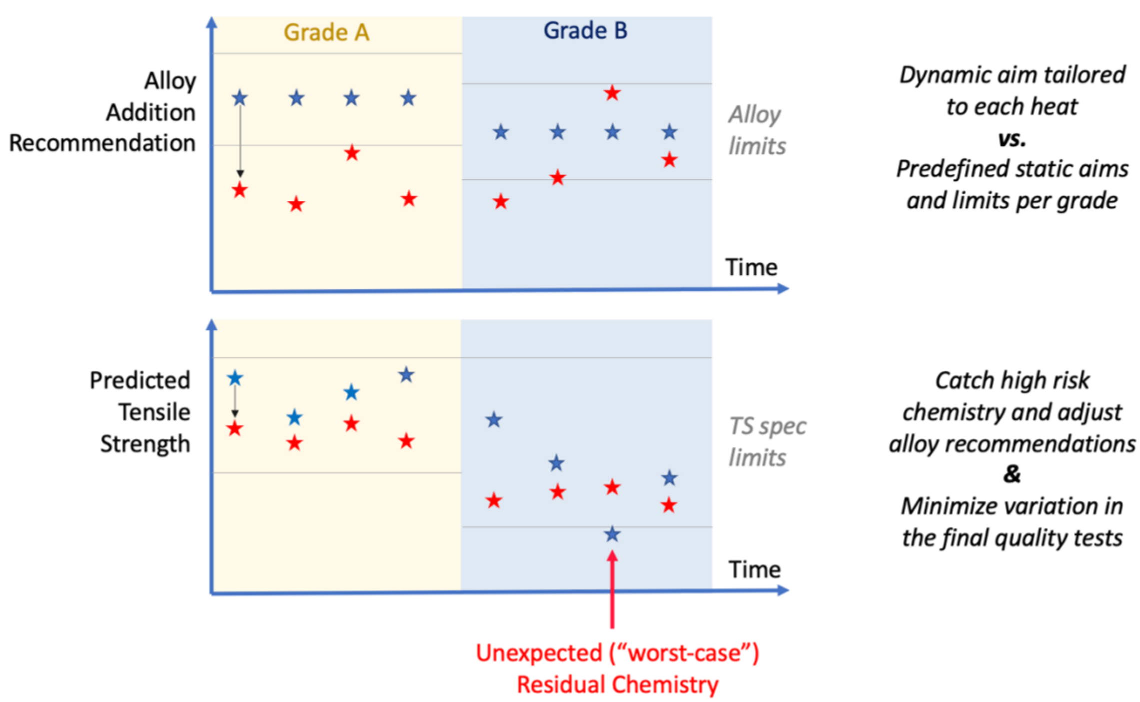

Figure 11: The ML software optimizes in real-time after each test sample, providing a personalized recommendation for each heat. While the software's alloy addition recommendation (in red) may differ from the standard recommendations (in blue), it leads to a reduction in the variance of the final test results, raw material costs, and production risk (see the RESULTS section).

Continuously improve data quality

Data quality is key to the success of any statistical analysis, especially in data-driven approaches like machine learning. Despite concerns about data quality, the implementation of machine learning tools is still achievable as long as the models are robust enough to handle any noise[7]. During this deployment the active use of white-box ML in operations has increased transparency in data and has streamlined the process of improving data quality. This created a positive feedback loop where the models highlight potential issues, the mills take action to improve the data, and in turn, the models become better.

The producer's mills utilizing ML in production have taken extra measures to improve data quality, such as replacing or fine-tuning sensors, upgrading data logging and compression processes, and refining testing practices. These practices are already part of operational improvement, but occasional problems can go unnoticed in the large amounts of daily production data. The presence of the online tool simply enables the mill to detect even the smallest problems and resolve them at a speed that was previously unattainable using conventional data analysis tools. The end-result is continuous improvement of data quality that underscores continuous improvement of production process.

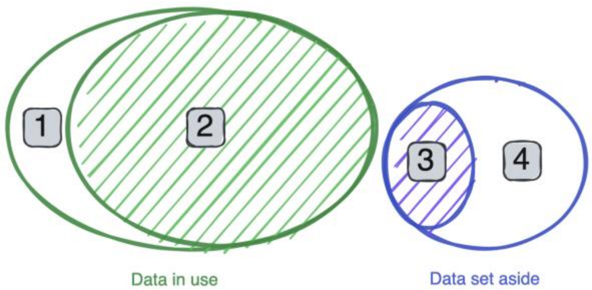

Figure 12: All test results are stored in the historian, but only a selected subset of the data is designated as valid for reference (green). This reference data generally includes the highest quality measurements (2), but it's important to note that this may not always be the case (1). On the other hand, data that is not marked as reference (purple), may still contain useful information such as healthy measurements (3) along with lower quality measurements (4). To ensure that only the best data is used, it is recommended to use tools that effectively processes the entire data set and allows filtering of high-quality data.

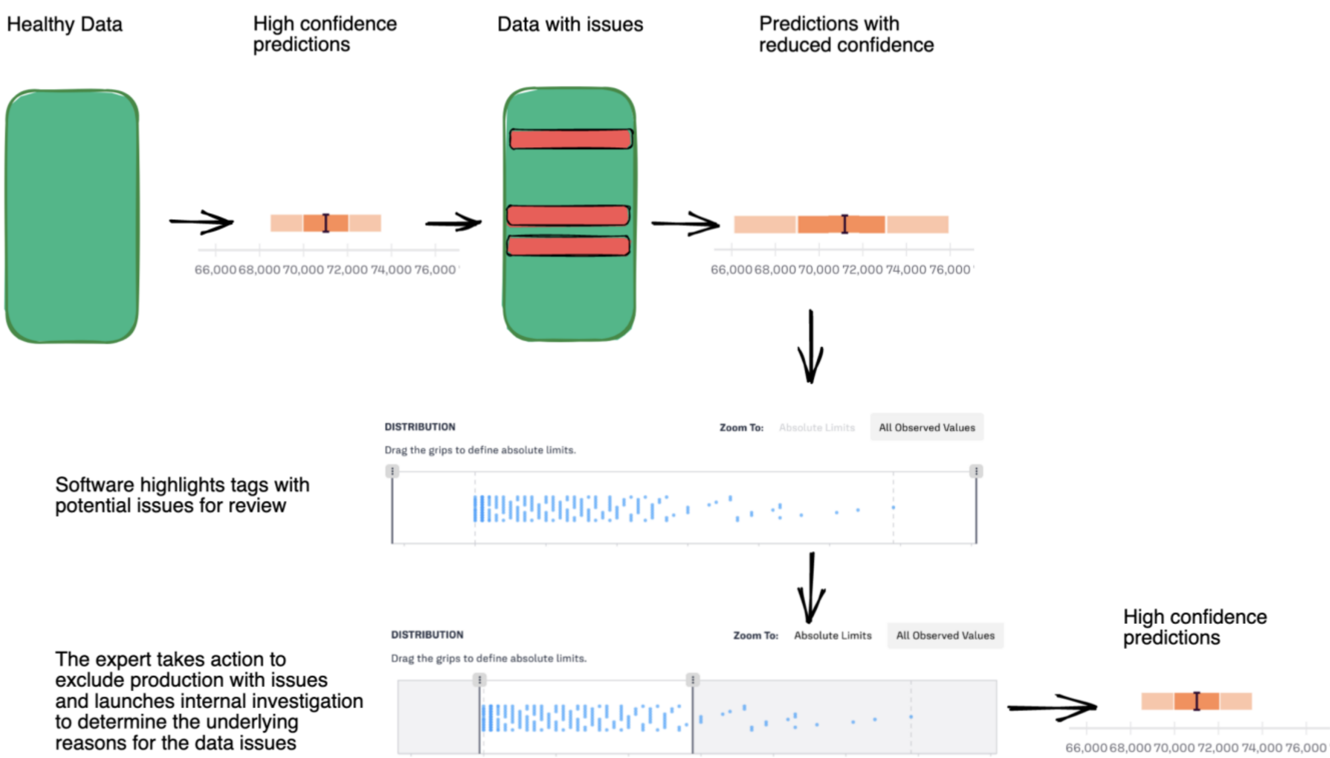

Figure 13: White-box models' prediction confidence will decrease if data issues such as faulty sensor readings or incorrect data entries occur. Since online models render predictions every new batch, this acts as a real-time alert to users. When the model's confidence drops for a specific period or product, the source of the problem can be visualized within the tool, allowing for further investigation at the mill to identify the root cause and address the issue in a timely and continuous fashion.

RESULTS

This section presents the results achieved by adopting ML based workflows to minimize alloy costs at the ladle furnace. The ML models accomplish this goal by rendering early prediction of final mechanical properties while the steel is still at the furnace and providing real-time recommendations to the operators. The operations and quality teams leverage the features and learnings emphasized in this paper daily to attain these outcomes. The two figures on this page show the long-term trends resulting from increased use of ML by the mill operations.

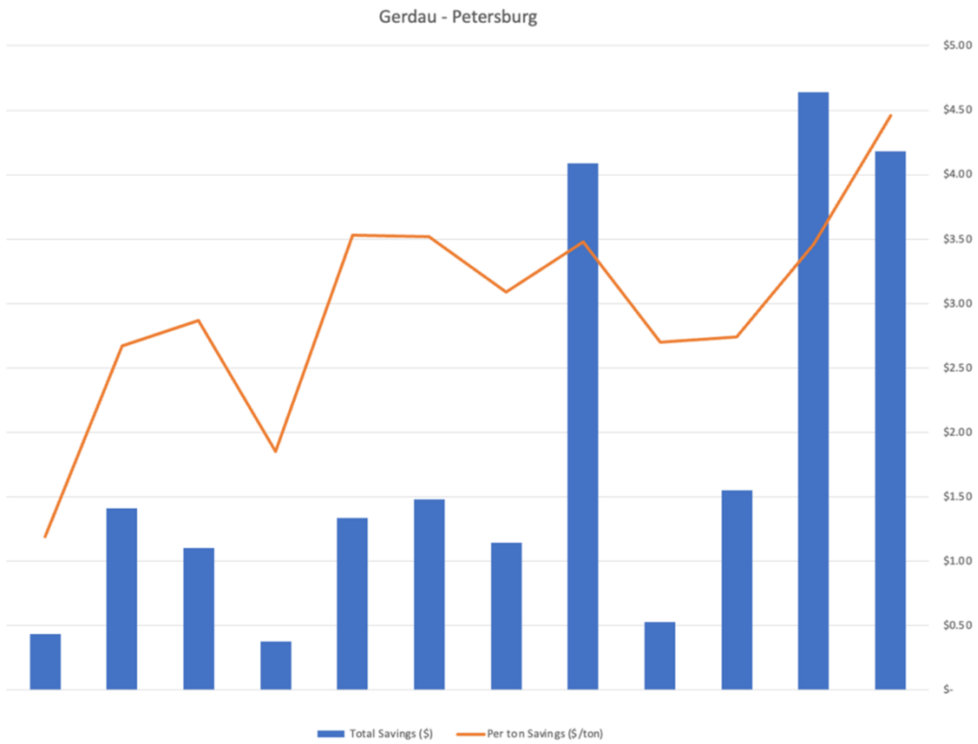

Figure 14: The annual savings cycle resulting from the operators using ML recommendations to optimize alloy additions at the ladle furnace. The blue bars represent the total savings achieved each month, while the orange line is the savings per ton of steel produced.

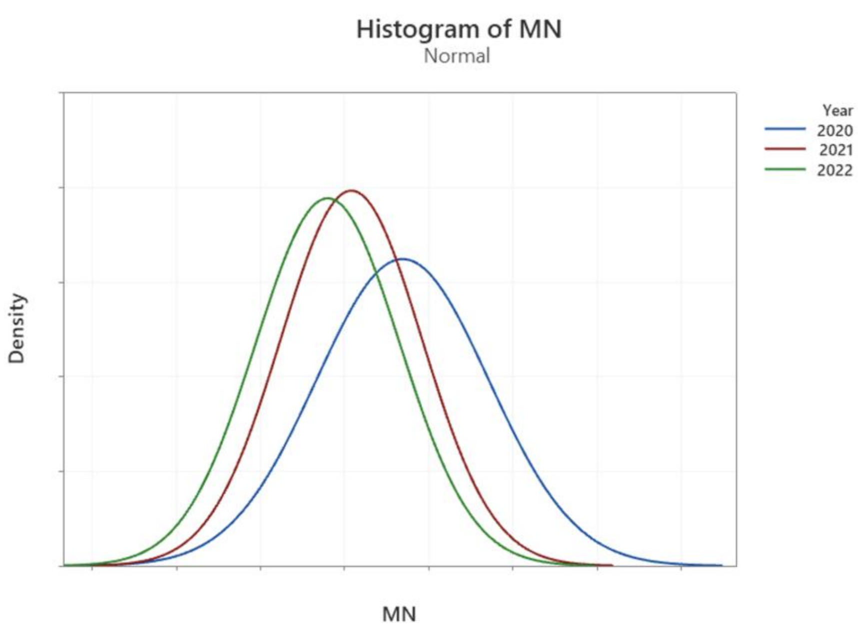

Figure 15: The distribution of manganese additions in popular grades has decreased over time due to the increasing efficiency of the operations team, who have incorporated key learnings listed in this paper. Furthermore, the ML models have also improved over time as more production data is fed into the software.

The following two figures examine the effect of utilizing ML recommendations on a chosen grade. These histograms highlight production in 2022 and illustrate the practical impact of the insights outlined in the "enhancing operational flexibility" section. The difference between the alloy content of the optimized heats of the grade is apparent in the first figure.

The second figure is noteworthy as it shows a roughly 15% decrease in the variability of the final quality tests in heats that use ML. This is a direct outcome of the operations team becoming more flexible. By adjusting the chemistry upstream on a heat-by-heat basis, the mill's operations are better equipped to compensate for the fluctuations in the incoming raw material and downstream rolling mill process. This proactive approach, based on predictions of the end-of-line test results, enhances the process capability while lowering costs.

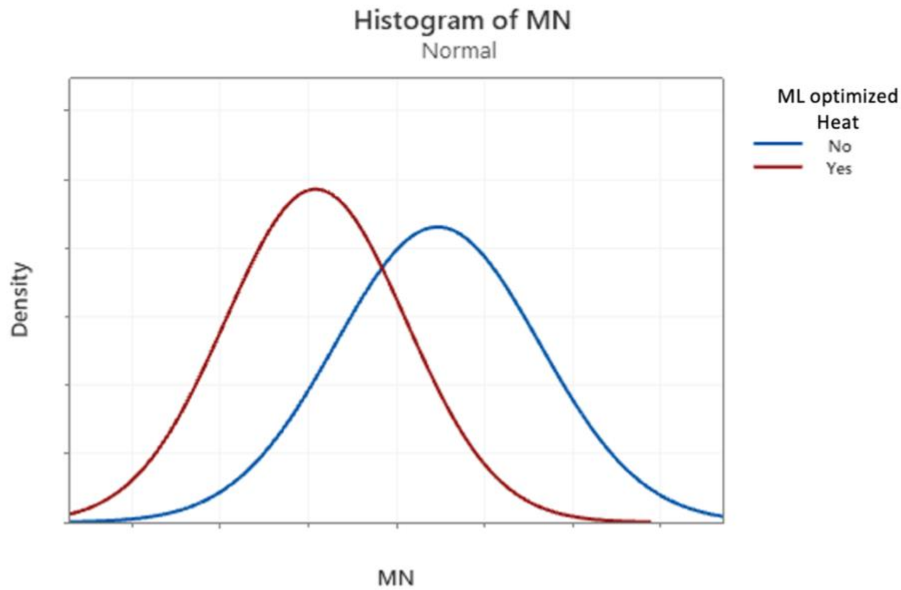

Figure 16: Comparing the manganese content of heats produced with and without ML recommendations in 2022 within the same grade. The priority at this stage of implementation is to ensure consistent production with ML optimized chemistry for the greatest financial gain.

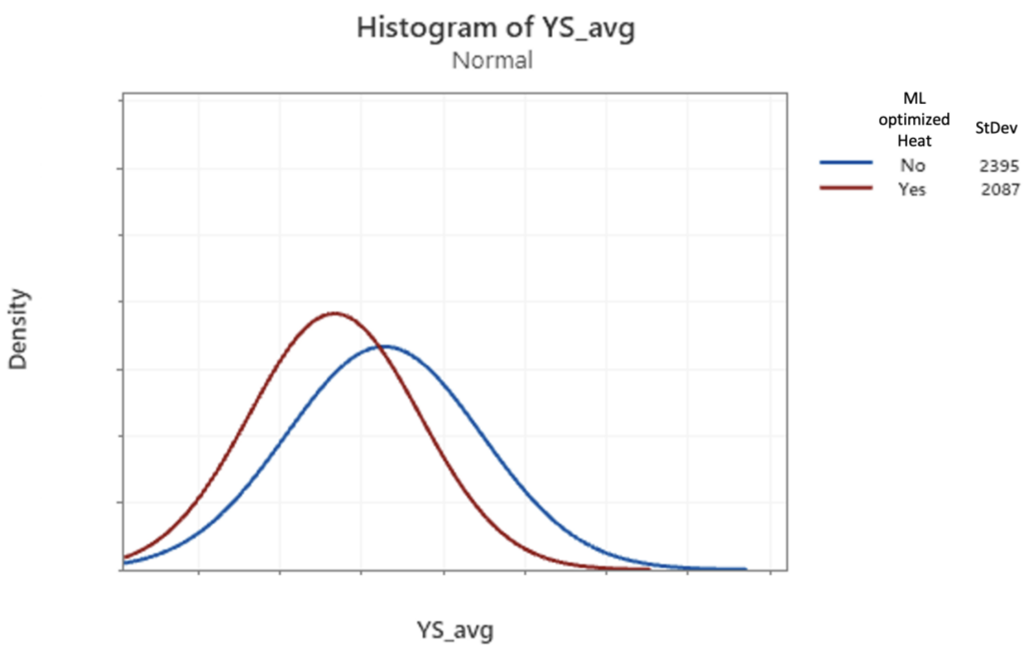

Figure 17: Comparing the distribution of yield strength in the products produced from heats that utilized ML recommendations versus those that adhered to grade book limits. There is a roughly 15% decrease in the yield strength variability in the optimized heats.

Building on the importance of flexibility, the plots on this page highlight the operational challenge of using ML for alloy optimization in the melt shop. An increase in the range of alloy addition is a natural consequence of following the recommendations of the ML model. When the model detects risk to final quality, it will suggest increasing alloys, and the opposite will occur if the risk is low. Adopting this new method of operation necessitates additional collaboration, training and coordination as discussed in this paper. The sooner these challenges are recognized and addressed during implementation, the quicker the production will be optimized with minimal risk.

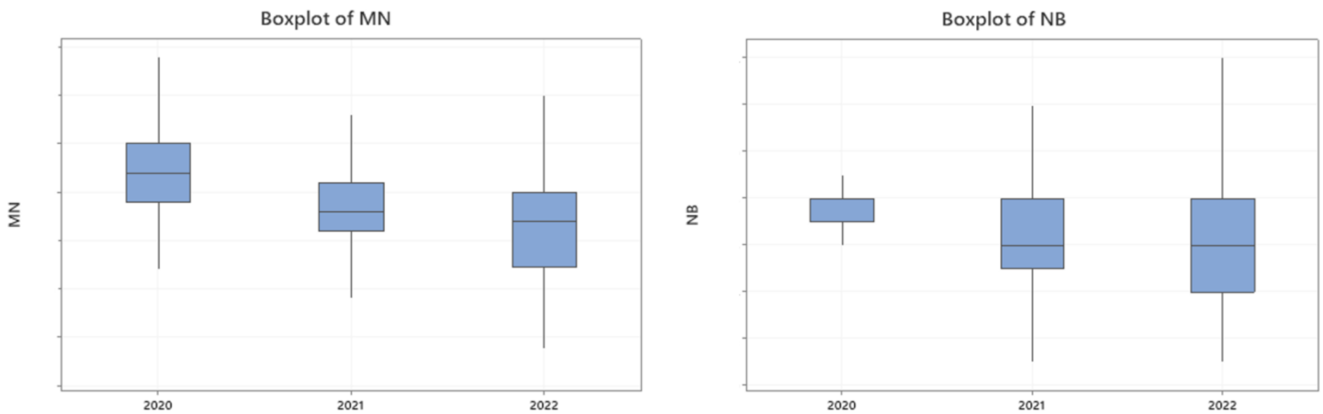

Figure 18: Optimizing steel chemistry in real-time at the melt shop requires new and more complex workflows for the operations team. The operations teams that can implement these additional steps benefit from decreased operation costs and enhanced process capability.

The operations team may not be ready for the workflow changes and flexibility required for real-time melt shop optimization. It's important to emphasize that although this paper highlights a compelling use case with tangible financial and environmental benefits, it is not the only practical application of ML software in the steel mill. Other potential uses for different stages of the production process can be quickly investigated by utilizing the automated software, as long as the data is stored and readily accessible. However, even in mills where the operational teams are eager to participate, the shift to a machine learning-driven workflow takes time as demonstrated by the results highlighted in this section. Those mills that start the transition earlier enjoy the benefits for a longer period and can better adopt to changes in the market and production conditions.

CONCLUSIONS

Leveraging machine learning to optimize production in real-time holds significant potential, but this potential is often unrealized due to projects getting bogged down in the early stages of ideation or modeling. In particular, a lot of effort is often put into the modeling phase without fully taking into account the implementation challenges or engaging the individuals who will be responsible for using these new tools on a daily basis. Vital aspects such as user friendliness, ease-of-maintenance and integration, and transparency are often overlooked in the pursuit of data and model accuracy

ML software can be applied to various applications as long as data and a well-defined problem exist. This paper focuses on its use in the melt shop, which has a proven and quantifiable track record across multiple mills. Moreover, the daily use of alloy optimization model across multiple divisions provides valuable insights into best change management strategies for scaling the real-world use of ML. Our hope is that the operational recommendations and the key software features highlighted in this paper act as a practical guide to other organizations as they embark on their digital transformation journey.

ACKNOWLEDGEMENTS

The authors extend their gratitude to the following individuals: Jonathan Fajardo, whose support with implementation and operator training proved invaluable; Martin McPherson, who established data connections to Fero and provided general IT support; Alice Pitchford and Carlo Travaglini, for their early support and efforts in driving this project forward; and John Boudreaux and Alp Kucukelbir, for their assistance throughout the implementation process. Their expertise and contributions were instrumental to the success of this deployment.

REFERENCES

[1]Sung Wen-Pei, Kuen-Suan Chen, Cheer Germ Go. "Analytical Method of Process Capability for Steel", *The International Journal of Advanced Manufacturing Technology*, *Heidelberg,* Volume 20, Issue 7, September 2002

[2]Andrei Paleyes, Raoul-Gabriel Urma, and Neil D. Lawrence, "Challenges in Deploying Machine Learning: a Survey of Case Studies", *ACM Computing Survey*, Volume 55, Issue 6, January 2022

[3] Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles C Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Burkner, Martin Modrak, "Bayesian workflow", *arXiv preprint*, arXiv:2011.01808, November 2020

[4]"Technology Trends Outlook 2022 Report", *McKinsey Global Institute,* August 2022.

[5]Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, Donald B Rubin, "Bayesian Data Analysis", *CRC Press*, 2013

[6]Alp Kucukelbir, Yan Wang, Bryan Williams, Taylor Giddens, and Jena Kreuzer. "Accelerating Product Chemistry Refinement Using Machine Learning", *Iron and Steel Technology Conference*, 2020

[7] Lukas Budach, Moritz Feuerpfeil, Nina Ihde, Andrea Nathansen, Nele Noack, Hendrik Patzlaff, Felix Naumann, Hazar Harmouch, "The Effects of Data Quality on Machine Learning Performance", *arXiv*, arXiv:2207.14529, November 2022