Machine Learning and Steel

Alp Kucukelbir, Fero Labs and Columbia University

Keywords: Machine Learning, Statistics, Continuous Improvement, Six Sigma

DRAFT

Lorem ipsum dolor sit amet, consectetur adipisicing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniam, quis nostrud exercitation ullamco laboris nisi ut aliquip ex ea commodo consequat. Duis aute irure dolor in reprehenderit in voluptate velit esse cillum dolore eu fugiat nulla pariatur. Excepteur sint occaecat cupidatat non proident, sunt in culpa qui officia deserunt mollit anim id est laborum.

A GENTLE INTRODUCTION TO MACHINE LEARNING

Machine Learning (ML) refers to a broad category of computational methods used for extracting patterns from data [1]. Unlike traditional computer programs, ML methods can identify patterns that are not explicitly programmed in their code. Engineers can leverage these patterns to identify causal relationships that affect production, explain past production issues, and optimize process parameters to reduce scrap and improve end-of-line quality.

To understand how ML works, consider the following core concepts. **Process:** a series of steps, often involving physical equipment, that produces a product. **Data:** historical or live measurements relating to a process, often taken from physical equipment. **Domain knowledge:** high-level information about a process, that need not be reflected in data. This knowledge is high-level in the sense that engineers who build and operate processes are clearly experts in their fields. But even the best engineer cannot immediately estimate mechanical or magnetic properties during production, taking into account the dozens (if not thousands) of factors that may affect it. A mathematical **model** of the process, however, is ideally suited for such a task.

ML uses historical data to build a mathematical model of a process (Figure 1). Consider the use case of optimizing mechanical properties. The mill produces data that can be stored in the format of a table. Each row corresponds to a product. Three columns contain yield strength, tensile strength, and elongation measurements; denote these as $\mathbf{y}$. The others contain alloy composition measurements, melt shop parameters, and rolling mill set points; write these as $\mathbf{x}$. ML can use this table to "learn" a multivariate function $\mathbf{y} = f(\mathbf{x})$. The model f is statistical in nature; it is extracted from data by analyzing high-dimensional correlations. A good model should represent known (and potentially unknown) physical relationships that drive the process.

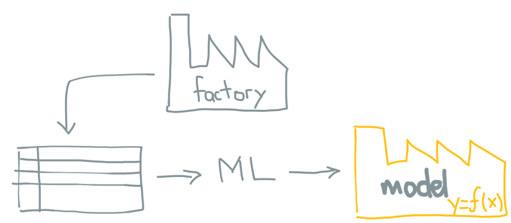

Figure 1: ML uses data to build a model of a process.

This highlights a critical assumption. The data should reflect the natural variation of the process. Measurements should contain common operating regimes and cover the majority of product types. ML strives to extract a function f that best describes the process; naturally, it is limited by the quality of the data provided.

What can an engineer do with a ML model of a mill? First, she can predict the mechanical properties of live production. This enables a workflow where melt shop operators can make cost-efficient decisions. Second, she can analyze and explain past production issues, such as "Why did a particular heat fail last week?" Finally, she can plan interventions. She can use the ML model to optimize a particular product to use less of an expensive alloy while meeting its mechanical property targets.

MACHINE LEARNING AND STEEL

ML is not new in steel. Since the 1990s, neural networks have been leveraged at the process control level [2]. What has changed since then?

Modern ML methods are now able to capture far more complex relationship. This increases the scope of its application. And recent developments in "white-box" ML steer away from the black-box nature of neural networks and drive towards presenting causal insights from data [3]. Figure 2 captures this transition.

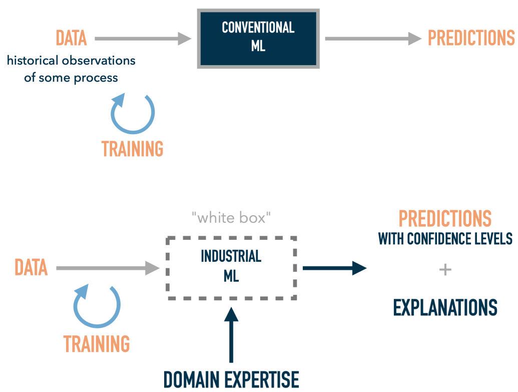

Figure 2: Modern Machine Learning.

The development and application of modern ML presents new requirements. How to integrate domain knowledge? How to build trust? These are the challenges of applying modern ML in practice.

MODERN APPLICATIONS IN PRACTICE

Alloy minimization in electric arc furnace (EAF) mills

EAF mills work by melting scrap metal. This presents a specific challenge: how to best adapt to the (unknown) variation of alloys and residuals contained within the scrap? Melt shop operations are designed to repeatedly measure the chemical composition of each heat and gradually add additional alloys. As a result, unnecessary amounts of expensive alloys are often consumed to ensure final products hit their targets.

Modern ML can help guide decisions at this stage. ML can calculate the optimal amount of alloy to add to a heat, while taking into account how the product will be rolled. Figure 3 visualizes this setup; note how the confidence intervals of each prediction matters. Melt shop operators that leverage these recommendations can obtain substantial savings in alloy consumption.

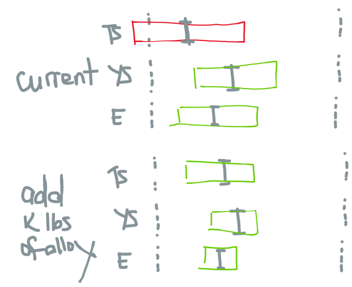

Figure 3: Alloy minimization.

Power loss minimization in electrical steel

Thousands of parameters affect the production of electrical steel. This presents a particular challenge when trying to diagnose the root cause of magnetic property deviations. With this many factors at play, a domain expert may struggle to narrow in on a set of factors that are worth studying. However, with modern ML, a domain expert can hone in on a set of causal insights that lead to a solution to the root cause of deviations.

DISCUSSION

How is ML different than statistics? ML and statistics are closely related. ML builds upon statistics to find complex relationships in data, quantify the uncertainty in measurements, and identify linear and non-linear patterns. ML methods are often designed to scale to big and messy datasets, as well as to live data streaming settings.

How is ML different than artificial intelligence? Artificial intelligence (AI) can be thought of as "decisions based on ML-driven patterns". For example, a self-driving car uses ML to detect stop signs in video data. An AI algorithm would then be the part that tells the car to stop in a certain way. Another common application area is robotics, where AI methods iteratively decide on how to move a robotic arm to accomplish a certain task.

REFERENCES

[1] Kevin P Murphy. *Machine Learning: a Probabilistic Perspective*. MIT Press, 2012.

[2] Martin Schlang et al. "Neural computation in steel industry". In: *European Control Conference (ECC)*. IEEE. 1999, pp. 2922–2927.

[3] J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.