Using machine learning to improve performance in engineering and beyond

A company can use the big data within its own organization to explain what truly drives engineering productivity, optimize resources, and predict the performance of teams.

Bill Wiseman, Chris Wigley, and Jacomo Corbo

For many companies, teaming and coordination are the missing links in their efforts to foster world-class performance. In engineering, this problem is particularly acute. Companies have applied lean engineering methods to raise value-added time and simplify processes, and they have introduced modularity and automation to increase the reuse of components and systems, and to replace repetitive tasks with automated processes. To maximize the benefits of these efforts, however, companies need to deploy effective teaming and coordination in the form of optimal staffing of project teams—an organizational structure that helps monitor their efficiency and promotes efficient work.

The opportunities to use training and coordination to improve engineering productivity are significant because many of its drivers are counterintuitive and not well managed today. One company's experience illustrates the extent of the opportunities.[1] Project starts and stops were among its main productivity killers: an 8 percent drop in productivity resulted from each one-week project stoppage. Overutilization also deserved greater attention—for every 10 percentage points of engineer utilization above 70 percent, productivity fell by 7 percent. The company also found that the size of a team and the geographic scope of its footprint matter. Every additional team member above seven and every additional time-zone region in a team's footprint reduced productivity by 7 and 5 percent, respectively. The company also determined that a sharper focus on group dynamics would probably improve productivity because it increased by 7 percent when the members of a team had previously worked together.

[1] All cited results are company specific. While the levers might be applicable more broadly, their impact will vary with the context.

A new approach developed by QuantumBlack, in collaboration with McKinsey, helps companies make engineering more productive through better teaming and coordination. A core element is QuantumBlack's proprietary data and analytics platform, Nerve, which uses latent data within the enterprise to derive insights. This tool extracts and learns the factors—such as team collaboration, the use of digital tools, and project management—that promote good outcomes across the project portfolios of organizations and makes it possible for them to apply these drivers to similar projects. For example, by using Nerve, a manufacturer of oil and gas equipment raised its productivity, on a project basis, by 14 to 37 percent as measured by costs, schedules, and quality. That equates to a 20 percent productivity lift across an organization with 7,000 engineers.

An analytics-driven approach

The Nerve tool analyzes broad ranges of disparate data that reside in organizations (but have rarely if ever been connected at scale before) to understand what actually drives engineering's productivity. The tool uses these insights to design and implement programs to transform it. In addition, companies can use Nerve to optimize resources (for example, supporting team-staffing decisions) and to predict project outcomes (say, by serving as an early-warning system).

Examples of analyses include the following:

- **Talent allocation.** Which engineers consistently improve the performance of teams?
- **Team composition.** Should teams that interact frequently be combined and located together? Should teams that do not interact at all be split up?
- **Cross-functional effectiveness.** Do reviews involving more than 20 people yield better results than a smaller setup would? Which people actually make a difference by participating?
- **Shape of pyramid.** What team structures lead to the lowest cost for delivering projects?
- **Automation effectiveness.** Are IT investments really improving productivity?

A three-phase journey

Companies deploy the Nerve tool in a three-phase journey led by a project team:

1. Inception. Typically requiring eight weeks, the inception phase entails ingesting multiple data sources and then cleaning, linking, and fusing them to create a "data lake." The project team also generates a list of hypotheses of key drivers to test. The legal and compliance frameworks for working with the data are put in place during this phase, as well.

Bringing together data sets from separate organizational silos that had never been linked previously can, by itself, provide an organization with valuable insights. At a carmaker, for example, the exercise of linking data exposed a remarkable failure of communication between the design and engineering teams. They typically spoke with each other only once a week while jointly developing a specific part.

2. Insights. The team starts by creating clusters of similar projects to compare "apples with apples" during the insights phase, which usually takes another eight weeks. Next, it uses machine learning to assess the drivers of time, cost, and quality and hunts for patterns across data sets. The team then creates an explanatory model to identify the drivers of performance and turns its attention to optimizing the allocation of resources—within real-world constraints, of course. At the end of this phase, the team develops visualizations of both the performance drivers and the output of the optimization.

The carmaker mentioned earlier, for example, captured fine-grained insights about when it should involve senior leaders in a development project. The company had assumed that their involvement would accelerate the project during all phases. However, the data revealed that while it does indeed have that effect during the early phases, it slows down the project and leads to rework during the middle and later phases. Similarly, involving suppliers was demonstrably helpful during the up-front and design phases but not during the others.

3. Impact. These diagnostic efforts come to fruition during the impact phase, which unfolds over the course of approximately 12 months, when the team applies the insights to design initiatives for change. To help the organization "live" the enhanced productivity, the team also embeds a version of Nerve in the organization's core systems as a predictive-intelligence tool to forecast future productivity and to design targeted interventions. As the machine learning algorithms consume more data, their performance improves. In addition, the team builds value-adding tools to assist

decision making. A tool that supports staffing recommendations, for example, can profile the kind of person who should work on a particular team and predict how that staffing choice would affect other projects.

One initiative the carmaker launched in response to the insights it gained from using the Nerve tool involved creating an elite concept-design team composed of both designers and engineers. This colocated team is responsible for kicking off the development process of each system in a vehicle and for coordinating the involvement of senior leaders and suppliers during the specific project phases when it is most beneficial.

* * *

Companies can apply the insights from the Nerve tool to improve productivity significantly by allocating their resources in a more effective way before making major organizational changes. The tool thus demonstrates that marginal changes, "when multiplied by a huge number of instances, or allowed to work over a long time, produce a significant effect."[2]

Bill Wiseman is a senior partner in McKinsey's Taipei office, Chris Wigley is a partner in McKinsey's London office, and Jacomo Corbo is a chief scientist of the QuantumBlack unit of McKinsey.

[2] "Schumpeter: Little things that mean a lot," The Economist, July 19, 2014, economist.com.