

Enterprise Learning

Curt Lefebvre, NeuCo Inc.

Artificial intelligence offers a revolutionary approach to asset optimization that is based on the behavior of the human brain.

OK, we all get it. Those aging assets that were obscured in the glare of traders' profit visions a mere two years ago are now the key to our very survival. Asset optimization is the current buzz phrase, and it seems like everyone is offering it. But what is asset optimization, and what do these offerings actually do?

Webster defines optimization as: "... an act, process, or methodology of making something (as a design, system, or decision) as fully perfect, functional, or effective as possible; specifically: the mathematical procedures ... involved in this."

To understand what someone means by asset optimization, you can start by asking yourself whether they are offering: 1) an act: the doing of a thing; 2) a process: a series of actions or operations conducing to an end; or 3) a methodology: a body of methods, rules, and postulates employed by a discipline.

There is no doubt that every power generation enterprise is inherently performing some degree of asset optimization. Every day, employees from the bottom to the top of the organization are making decisions and performing acts intended to make their assets fully perfect, functional, and effective. There are also plenty of engineering and consulting firms that can offer specialized acts, or train your employees with methodologies designed to improve the effectiveness of their acts.

I'm sure that we can readily agree that not all of these acts have the desired effect, and even the ones that lead to a positive impact still leave significant room for improvement. It's not easy in large organizations for employees to understand the impact of the acts they perform, much less know the optimal combination of acts they should be performing.

What if it were possible for every individual in a power generation enterprise to better understand the business impact of key decisions they make, prior to making them? What if it were possible to implement a real-time process for continuously improving the performance of your assets? What if you could coordinate key local decisions or actions toward the global impact they have on the company's bottom line?

This sounds like a pretty tall order, and it is. Power generation enterprises are by definition massive, complex, distributed, and highly dynamic. They are driven by multiple and often

conflicting objectives and must operate within tight constraints. Hundreds of thousands of decisions across equipment, units, plants, and portfolios are executed within an enterprise each day

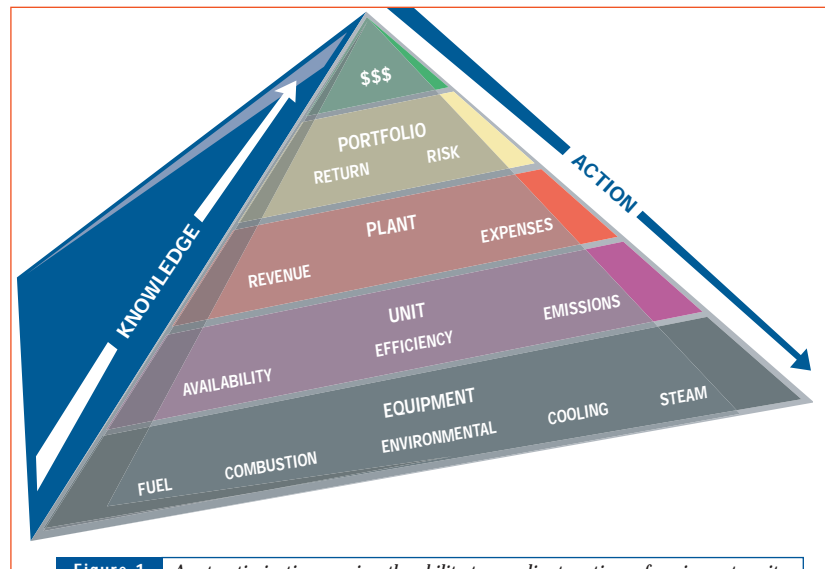


Figure 1 Asset optimization requires the ability to coordinate actions of equipment, units, plants, and portfolios in a dynamic market toward a single business goal.

(see Figure 1), so is it realistic to expect that we can capture all of this data, integrate it, and then convert it into coordinated action toward a single objective?

One thing is for sure: putting monitoring tools in the hands of operators and digital dashboards in the hands of executives will come nowhere close to solving the problem. Monitors and dashboards only tell you what is happening; much more important is knowing why it is happening and how to change it. It is critical to get information about what is happening in the hands of executives so they can make better strategic decisions. It is even more important to get knowledge of why it is happening and how to change it in the hands of those operating the plants, the people who have a real-time impact.

Curt Lefebvre is president and CEO of NeuCo Inc., a leading provider of optimization software solutions for the electric power industry. Prior to NeuCo, Mr. Lefebvre founded NeuroDimension Inc. to help bring advanced forms of artificial intelligence technology from the laboratory to the marketplace. He is an author of the academic textbook *Neural and Adaptive Systems*, and he holds a Ph.D., M.S., and B.S. in electrical engineering.

But how is it possible for control room operators to understand the repercussions of their local actions – pulling a lever here or there – on the company's bottom line? How can maintenance engineers plan maintenance activities to maximize not just plant up-time but also overall profitability? How can units be dispatched in a way that accounts for current variable costs and future availability in the face of highly uncertain power, fuel, and allowance markets? How can all of these decision makers, acting over vastly different time scales, adapt as plant and market conditions change?

The answer to the problem is a revolutionary approach based on our understanding of the mechanics of the human brain, known as learning theory. This artificial intelligence-based approach implies a radical shift in mindset from deductive to inductive reasoning and opens up an entirely new world of asset optimization possibilities and opportunities.

Incredible Machine

The brain is truly the most incredible machine that we know. Just like an enterprise, it is massive, complex, distributed, and dynamic. Unlike today's typical enterprise, however, it has billions of parts that work together with an amazing degree of coordination. The workers in the brain continuously cooperate to solve global problems inductively and with relative ease.

To explain why a learning theory approach offers such promise for optimizing power generation enterprises, it is first necessary to describe how the brain functions (see Figure 2).

The brain is a collection of approximately 10 billion decision makers called neurons. While these neurons are massively interconnected, each neuron is only synaptically connected to a few thousand other neurons. This means that very few neurons in the brain have a direct connection to the external world. They can only communicate with the local neurons that surround them.

What we know about neurons is that they behave in a binary fashion. They either fire or they don't fire in response to other neurons that fire in their vicinity. When they do fire it is because they have received a precise pattern of input activity from neighboring neurons that happened to fire at the same time. This leads to a cascade of firing neurons. Those at the beginning might have been responding to visual stimulation from seeing a ball approach and those at the end could be driving motor movement in the arm to allow us to catch it.

Credit Assignment

We know that the input patterns a neuron responds to have been learned from correlated past events, and that the neuron continues to adapt its response based on the global success of the task at hand. The concept that is at the heart of learning theory is known as credit assignment. How does a neuron deep within the recesses of the brain, which has no direct ability to see whether or not the ball was actually caught, get credited or penalized for its role in performing the task?

There are lots of theories and algorithms for the credit assignment problem, some more biologically plausible than others, but what is significant here is that they all provide a process, specifically the mathematical procedures involved in this, for coordinating a massive

This artificial intelligence-based approach implies a radical shift in mindset from deductive to inductive reasoning and opens up an entirely new world of asset optimization possibilities and opportunities.

number of local decisions toward global objectives.

Learning theory assigns credit to the neurons that fire in response to a particular situation, and this is a critical underpinning of this proposed new approach to asset optimization. What if the same kind of credit could be applied to each of the thousands of local decisions taken each day within a power generation enter-

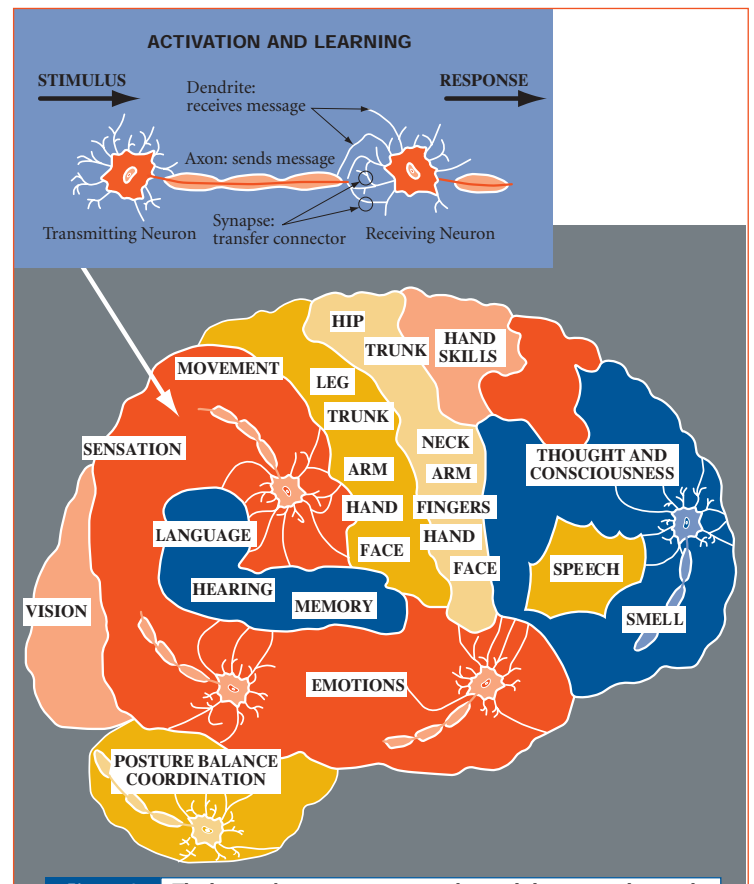


Figure 2 The human brain is massive, complex, and dynamic and can solve global problems inductively and with relative ease.



Figure 3 *Enterprise learning makes it possible for a power generation enterprise to adapt like a living, breathing, learning being.*

prise? With that kind of “enterprise learning” capability, we would have the ability to adapt and adjust local parameters to affect a desired global outcome. A power generation enterprise could be likened to a living, breathing, learning being.

Leveraging Investments

This may seem like a futuristic dream, but the technologies exist today for making it a reality. Furthermore, power generators have already made most of the required investments in instrumentation, automation, and monitoring. These investments are just waiting to be tapped. We have already started leveraging these investments on a much smaller scale using artificial neural networks to solve local optimization problems. Ten to 15 percent of US generators have real-time combustion optimization systems and

WEblink

For
MRI use
only.

together are saving an estimated \$250 million per year from the NO_x and heat-rate reductions they provide. Similar optimization systems have recently been implemented on other subsystems as well (e.g., soot blowing, steam temperature, selective catalytic reduction, selective non-catalytic reduction, and flue gas desulfurization systems).

The current systems are only scratching the surface of the potential of learning theory, however. People are approaching the optimization problems independently, building small brains to solve each problem without any synaptic connections between them. This will not work. The systems we are trying to optimize are highly interdependent, and it will not be possible, and could even be dangerous, to optimize them without integration. Learning theory provides the ability to perform this integration, but it will require a radically new approach to the design of the software systems from which these optimizers are built.

Fortunately, there is such an approach to software design now available. Microsoft and others have invested billions in shifting their focus from PC to enterprise applications. Bill Gates has bet the company on a radically new framework for developing all software applications, called .NET. The result provides exciting opportunities for innovation across the board, but in particular to the application of learning theory to the challenge of asset optimization. I refer to this marriage as enterprise learning.

This concept is beginning to gain momentum. In fact, the US Department of Energy recently partnered with NeuCo to invest \$18 million in a four-year demonstration project through its Clean Coal Power Initiative to prove out the potential of integrated optimization systems within the power industry. The initiative will leverage and accelerate the investment NeuCo has already made in the commercialization of a suite of optimization products in a real-time integrated software environment. The result will be one large integrated brain that provides real-time optimization of the equipment, units, and plants in a power generator’s portfolio toward business objectives.

Summary

Asset optimization requires more than monitors in the control room and digital dashboards in the boardroom. It requires arming the people running the plants with the ability to translate business goals into operational actions. A new form of asset optimization achieved through the application of learning theory makes this possible and it will have a tremendous impact on the future of power generation enterprises. The technology to build truly adaptive enterprises will be available to progressive power generators within five years, and it will be deployed throughout the industry within 10 years. The power generation industry, with its heavy investments in data, control, and automation systems, is able to take the benefits of this disruptive technology to a whole new level – to build truly adaptive, learning enterprises. ■