As seen in the February 2022 issue of **Nuclear**News


# Powering our nuclear fleet

# with artificial intelligence

By J. Thomas Gruenwald, Jonathan Nistor, and James Tusar

## Hidden treasure

We've all heard the stories of lost treasures being found in dust-filled attics, locked away in forgotten wall safes, or hidden in secret compartments of antique desks. Some of these true accounts, such as a rare copy of the Declaration of Independence hidden behind wallpaper or an authentic Van Gogh relegated to collecting dust in an attic, can lead to seven- and eight-figure jackpots when the discoveries are made.

What about our own treasures locked away in long-forgotten data storage drives or plant process computers? Imagine that you could gain keen insight into every operational issue you have by using the data you've been collecting for decades. In a nuclear power plant, data is routinely generated and collected for a myriad of purposes—whether it be for core monitoring, exposure accounting, equipment monitoring, or other reasons. While that data may serve its primary function exceedingly well, the information contained within it and in the aggregate is profoundly richer than most could imagine.

Worldwide, over 40 percent of companies have leveraged their data to some degree to enjoy a diverse set of benefits. Topping that list are better understanding of customer behavior, improved control of operational processes, better strategic decisions, and cost reductions. Furthermore, those organizations that are able to quantify the gains from leveraging their data have reported an average 8 percent increase in revenues and a 10 percent reduction in costs.[1] While the nuclear energy sector may be late to the party, there is still time to reap the value hidden within its data before it is downsampled for archiving, corrupted beyond repair, or lost entirely.

*Continued*

1 C. Bange, T. Grosser, and N. Janoschek, "Big Data Use Cases 2015," BARC Research, July 2015; barc-research.com/research/big-data-use-cases-2015.

## Current industry landscape (a mixed bag)

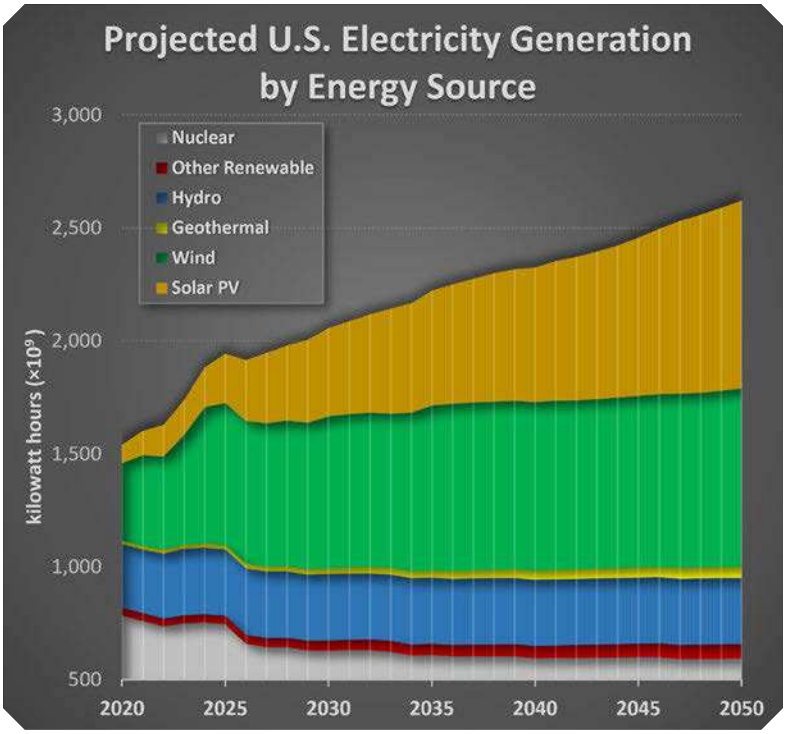


The outlook for the existing nuclear fleet is a complicated one. On the one hand, the need for a robust nuclear energy sector has never been greater. The overarching need to transition to a carbon-neutral energy landscape should place nuclear as the favorite to replace coal, natural gas, and oil as the baseline energy source. The electric grid requires reliable "always on" electric power that can be supplemented with other forms of low-carbon-emitting energy. While solar and wind are available 24.5 percent and 34.8 percent of the time, respectively, the nuclear backbone is available 93.5 percent of the time in the United States.[2]

In one sense, the nuclear industry is experiencing a resurgence. Public opinion and political will have followed the drive to limit or eliminate carbon emissions. People generally see nuclear as a safe, carbon-free solution to preserve the environment. These shifting winds are signified by the state of Illinois recently recognizing nuclear energy as critical to achieving its clean air goals and committing more than $700 million in support of Constellation's nuclear power plants to avoid several premature plant shutdowns.[3]

On the other hand, while the nuclear industry has been providing clean, safe power for over 60 years with a carbon-free footprint and a pristine safety record, making it a keystone of U.S. carbon-free energy production, there remain economic forces that challenge its long-term viability. The industry's unique regulatory environment and higher generation costs than those associated with fossil fuel plants are both factors that sit squarely at odds with our climate objectives.

Consequently, while many countries are bringing new nuclear plants on line to provide carbon-free energy, the share of power generation from the domestic nuclear fleet is shrinking. As of May 2021, 52 reactors were under construction worldwide, with China planning to build 150 new reactors over the next 15 years. Compare this with only two reactors under construction in the United States!

The energy outlook for renewables is a favorable one, with the share of power generation from wind and solar doubling between now and 2050.[4] This is only part of the equation, however. Nuclear energy will be vital to meeting our collective climate goals with a clean energy mix of wind, solar, and nuclear. To ensure the continued viability of nuclear energy, we must deepen our understanding of key aspects of nuclear power generation and strive to continuously reduce generation costs using the latest, most effective technologies available.

There is a clear opportunity to apply artificial intelligence (AI) and machine learning (ML) to improve nuclear plant efficiency and reduce costs. AI can be used over a wide range of nuclear plant operations, from predicting component lifetimes and evaluating asset health to understanding core dynamics for more accurate reload planning and economical fuel purchasing. The application of AI/ML to reload core design has been a key player in reducing reload fuel costs, which account for 20 percent or more of total power generation costs.

2 "What is Generation Capacity?" U.S. Department of Energy, Office of Nuclear Energy, May 1, 2020; energy.gov/ne/articles/what-generation-capacity.

3 "Bill to preserve Illinois nuclear passes legislature," *Nuclear Newswire*; ans.org/news/article-3247.

4 Annual Energy Outlook 2021, U.S. Energy Information Administration; eia.gov/outlooks/aeo.

## The reload process

Today's operating nuclear plants go through a complex process of designing the reactor core as part of the reload process (i.e., preparing to reload the core for a new fuel cycle). There are many design constraints, goals, and limits that must be met during the reload process to ensure a safe and economically viable fuel cycle, such as total energy production, fuel exposures (for fuel pellets, fuel rods, and fuel bundles), radial bundle power, and thermal limits.

In a typical 24-month nuclear fuel cycle, approximately one-third of the fuel is fresh, one-third has been in one prior cycle, and one-third has been in the reactor for two fuel cycles. The cost of fuel for a two-year fuel cycle can be as high as $100 million. Thus, the reload process employed by nuclear fuel departments is highly leveraged economically, and improvements can yield significant cost savings. Moreover, the reload process spans approximately one year from establishing the design goals and limitations, designing the fuel and reactor core, licensing the core, fabricating the fuel, shipping the fuel to the nuclear station, loading the new fuel into the core, and preparing the startup and operating plan for the two-year fuel cycle. The energy obtained from each fuel bundle must span three fuel cycles and deliver the planned energy over its six-year lifetime. If the exploitable energy is somehow miscalculated, there is an impact on the energy output for the next six years, and thus the economic performance of the reactor may be materially significant.

Calculations in the reload process, which predict energy output six years in advance, are daunting and complex tasks. The planning team must navigate a labyrinth of regulations that apply to the complex components of the reactor while accounting for a myriad of nuances in the fuel that impact ultimate performance and the ability to meet thermal limits while delivering rated power.

Fuel is not just fuel. The distributions of uranium, enrichment, and burnable absorbers such as gadolinium all vary throughout the array of assemblies. The distribution is aimed at meeting safety and operational requirements while minimizing fuel costs. In addition, in light water reactors, the water, which acts as both a coolant and a moderator, flows through channels in the fuel assembly. To complicate matters further, new fuel designs are introduced periodically. New designs not only change the mechanical design and composition of the fuel but also modify the way coolant flows through the bundle, which has a ripple effect on energy output.

## Stubborn reload design problems

Historically, several problems have persisted that affect the ability to further improve the economics of reload fuel planning. These problems can limit reductions in the amount of fresh fuel required to be loaded into the core (known as the reload batch size), resulting in excess direct fuel costs. In addition, they can have an impact of power generation if the core is less reactive than expected, or by the potential need to derate power if conditions require it.

### *Key problems*

■ **Inability to predict moisture carryover**—The amount of moisture mixed with steam leaving the reactor's moisture separators, referred to as moisture carryover (MCO), has been nearly impossible to predict by conventional methods. There are design specifications limiting

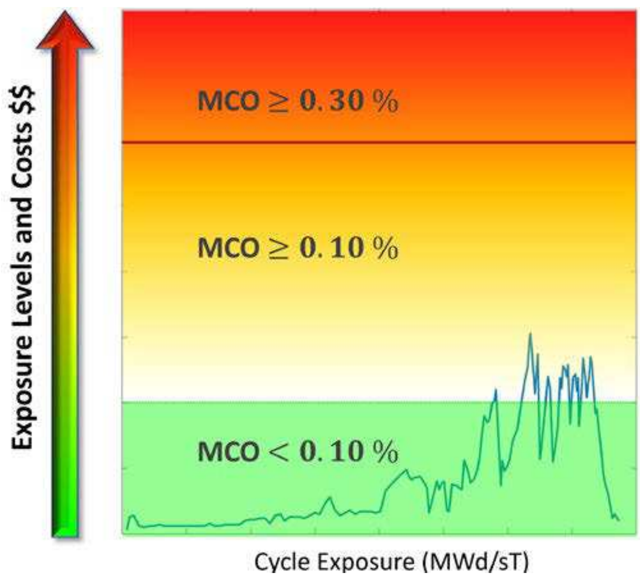


Relationship of high MCO to cumulative radiation exposure and costs at a given generating station. High MCO can lead to accelerated rate of erosion of main turbine components (above 0.10%) and accelerated rate of erosion of main steam isolation valve internal surfaces (at 0.30%).

*Continued*

how much MCO is permissible before the operator must take remedial action (of which one costly option is a power derate). Excess moisture in the steam is problematic for many reasons, most importantly due to the ability to carry impurities dissolved in the water throughout the entire plant. MCO can increase erosion of the internal surfaces of the main steam isolation valves and at the turbine, potentially causing costly repairs. Perhaps even more troublesome, soluble cobalt-60 is carried over with the steam, increasing plant dose rates and the collective radiation exposure of plant personnel. Beyond this, a small reduction in electrical output is expected with high MCO. Until recently, the primary method to mitigate high MCO was to design the core with a larger-than-required reload batch size, thereby introducing potentially unnecessarily high reload fuel costs.

■ **Unpredictability of eigenvalue in BWRs**—The hot reactivity parameter of the core (known as $k_{\text{effective}}$ or simply the eigenvalue) is one of the most fundamental parameters in nuclear engineering and has been notoriously difficult to predict accurately in boiling water reactors. Its trend directly affects the energy capability of the reload core, and an inaccurate eigenvalue projection can be costly. If the actual eigenvalue is higher than predicted at the rated power, then the designed core's reactivity is less than expected, thus leading to less generation output than desired (costing perhaps upward of $1 million per fuel cycle). Whereas if the actual eigenvalue is lower than predicted at rated power, then the designed core's reactivity is greater than necessary, and more fuel was purchased and loaded than required (potentially costing upward of $1 million more than necessary).

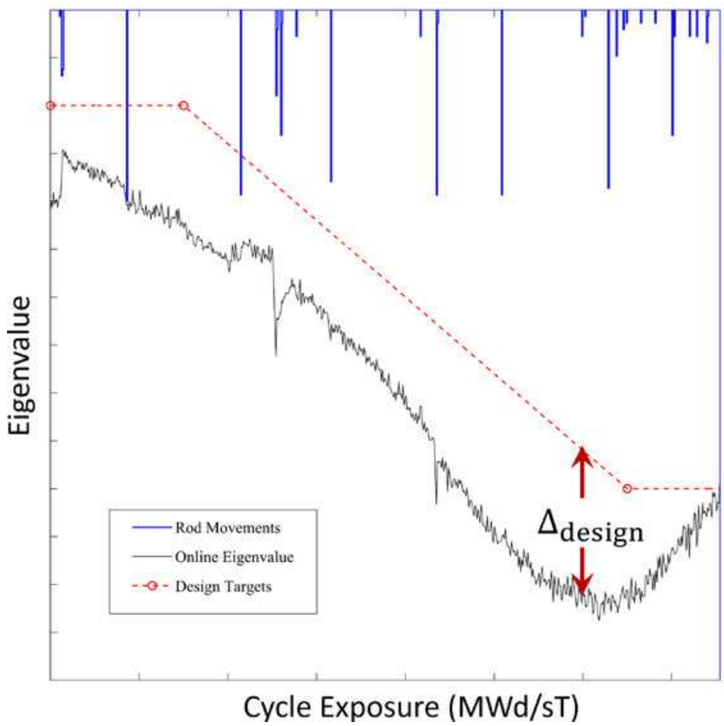

Disparity between the on-line eigenvalue and design targets obtained through conventional means.

Conventionally, eigenvalue predictions rely on estimates made by a committee of experienced core designers looking at past eigenvalue behavior and the characteristics of the reload core being designed. This approach has its limitations, especially when new fuel or core designs are introduced, and on average has been sufficient to achieve a deviation $\Delta \sim \pm 0.002$ between the design and on-line eigenvalue. The possibility exists to reduce this deviation fourfold, thereby leading, potentially, to millions in annual savings.

■ **Uncertainty in predicting on-line thermal limits**—Compliance with technical specifications for thermal limits is essential for operating reactors. Core designs include margin to these limits to prevent challenges to the operators. Core designs are performed with what is called "off-line" nuclear methods, which have no feedback from in-core nuclear instrumentation, while actual core monitoring is performed with "on-line" nuclear methods, which *do* have feedback from the in-core nuclear instrumentation. These differences cause a bias between the off-line and on-line thermal limits, which must be taken into account during core design. This is a challenge, since the bias varies from cycle to cycle, which makes its quantification difficult. The same principle described above applies, where too much margin results in increased fuel costs and less-than-adequate margin results in operational challenges and potential decreased generation revenue.

Accurately predicting the behavior of these important attributes has been extremely challenging for BWRs, because the dominant mode of operation in a BWR is a complex two-phase flow in the upper part of the core that in turn affects the reactivity. Incomplete understanding of the physics of two-phase flow in this region results in a costly degree of uncertainty. These problems limit the exploration and realization of more economical fuel loading strategies that, if solved, could lead to a 10 percent reduction in fuel costs in the aggregate. Importantly, AI and ML show significant promise for solving these problems, with solutions already being deployed in many BWRs across the domestic fleet.

## Energizing reload design with AI

In 2017, Blue Wave AI Labs and Constellation (formerly Exelon Generation) began working together to apply techniques rooted in artificial intelligence to solve some of these problems. Constellation operates the largest fleet of nuclear plants in the United States (14 BWRs at eight generating stations, and 7 PWRs at four stations) and, consequently, has a substantial amount of design, performance, and operational data relevant to these key challenges. In particular, the MCO and eigenvalue predictability problems were the first to be tackled, and solutions to both are being integrated into Constellation's reload design process.

### *Machine learning for nuclear power*

First, a little jargon. Machine learning is a branch of artificial intelligence that extracts answers to complex problems that may be intractable by more conventional means. It is especially useful where large amounts of data are available corresponding to unusually complex nonlinear problems that aren't solvable by analytical techniques or physics-based models. While there are many branches within ML, the objectives broadly fall into several categories:

- *Supervised learning*
  - *Regression*—used to predict a continuous variable such as thermal limits, MCO, or $k_{\text{effective}}$ at a given reactor statepoint.
  - *Classification*—used to assign elements to one of many categories (for example, equipment monitoring through determination of diagnostic health states).
- *Unsupervised learning*—used for clustering, anomaly detection, dimensionality reduction, and feature engineering.

The underpinning of ML is the universal approximation theorem, which guarantees that an artificial neural network can represent a true function, $F(x)$, to an arbitrary degree of accuracy if certain straightforward conditions are met. Even more important is the existence of sufficient data with a distribution that approximates the distribution expected for the target system.

The training data is the dataset used to infer this function and consists of many historical observations. The function's inputs, or raw features, are the statepoints of the reactor $\vec{x}_i$ at a given instant in time, whereas the targets are the corresponding outputs of the true function, $y_i = F(\vec{x}_i)$. The MCO measurements, on-line eigenvalue, or thermal limits are the training targets for the three problems described.

Fundamentally, a reactor statepoint is the collection of all the information necessary to completely describe the state of the core at a given instant. In practice, we must rely on limited information that is known through measurement, design, set point, or simulation. For example, the measurement of operational parameters such as thermal power and core flow, the control rod pattern plus notching, fuel and lattice designs, and the plethora of outputs from the core simulator collectively form an approximate representation of the core. By having enough observations $(\vec{x}_i, y_i)$, the underlying function that governs a process can be learned.

### *It's all about the data*

Returning to the problems at hand, each two-year fuel cycle contains hundreds of daily reactor statepoints. While each cycle may contain hundreds of points, in one respect, the fuel cycle itself can be regarded as a solemn point that codifies all the information pertaining to that designed core. As such, it is crucial for data from multiple fuel cycles to be pooled together in the training set to learn the complete functional dynamics.

For multi-reactor sites, it may be possible to combine the data from each unit if the underlying function is expected to be similar. Constellation has between six and eight fuel cycles worth of data across most of its BWRs. While this may seem like a lot (tens of thousands of datapoints in the aggregate), typical applications of ML, such as image recognition, require millions of training samples. A number of techniques have been employed to enhance the datasets, including data augmentation for maintaining expected distributions, interpolation of training targets, and transfer learning to take maximum advantage of information from multiple sites. These techniques have made it possible to extend the development of highly accurate models to reactors possessing less data than would otherwise be required.

Another challenge to overcome concerns the decision of what input features are necessary for training. Oftentimes this becomes the major obstacle in adopting AI within a

*Continued*

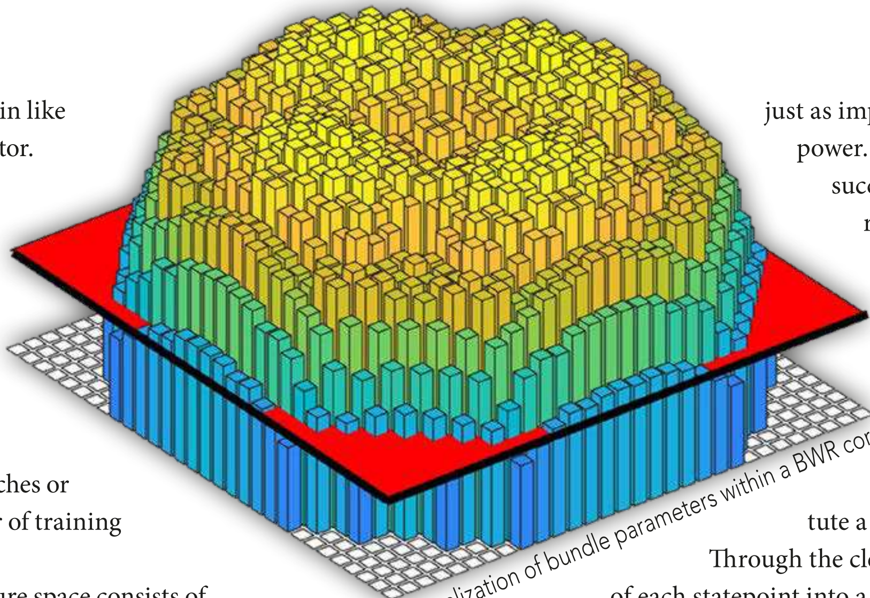

Visualization of bundle parameters within a BWR core.

very specialized domain like the nuclear energy sector. A pure data scientist may approach the problem with the attitude "the more inputs the better," but that perspective does not work when the sheer number of inputs matches or outweighs the number of training examples.

Here, the input feature space consists of tens of thousands of bundle and nodal outputs from a core simulator, hundreds of thousands of pin-by-pin fuel attributes, and dozens of global reactor variables. Left unconstrained, all this information can be used to train a model with precisely zero training error and absolutely zero predictive power! Where's the utility in that?

Too many variables and too much training cause the model to be useful for the training set only, not more general situations. When this occurs, the model has fit the noise in the data, masking the underlying functional dynamics, and the model is said to be overfit (comically so, in this example). Likewise, the complexity of the model's architecture (the number of neurons, for example) also contributes to the likelihood of overfitting. Conversely, when a model continues to perform well in new situations (e.g., new fuel cycles), it is said to "generalize" well. Generalization means that the underlying dynamics governing the process are captured well by the model, and the training process is stopped before latching onto the random noise within the data.

The trick is to find a balance between (1) the size and nature of the input feature space and (2) the model architecture—collectively, the modeling methodology—and to meticulously validate the methodology in order to arrive at the most generalizable model. For MCO, the answer is to reduce the input feature space to a "canonical set" of key drivers of MCO through feature engineering and a physical understanding of the underlying mechanism. In doing this, the input feature space is reduced to a few dozen key variables that capture the dynamics of MCO. This allows us to develop models with parameters that operators can control, giving the models not only predictive power but, just as important, corrective power. For eigenvalue, the successful approach relies more on the nature of the model's architecture, while retaining the vast collection of input features that constitute a reactor statepoint. Through the clever transformation of each statepoint into a three-dimensional image of the reactor core—viewed through various "filters" of exposure, void, power, and so on—we exploit a convolutional neural network architecture, which has been shown to be very effective with tasks like image recognition and natural language processing.

### *Predictive power*

As with any innovation, results are the ultimate arbiter of its value or utility. For MCO, an example of this predictive capability is illustrated in the graphs on the next page, where the model predictions are stacked up against MCO measurements for two of Constellation's units. Here, the model predictions are obtained from the exposure accounting collected throughout the cycle.

Since the time the model was first deployed, and over the past three years, the average prediction error is ±0.018 percent MCO at this station. This exceptional level of performance is now limited only by the resolution imposed from the MCO measurement uncertainty. Similar levels of accuracy have been obtained at the 10 additional BWRs that have adopted this enabling technology.

Also shown in the graphs, the eigenvalue model performance demonstrates a fourfold reduction in prediction uncertainty when compared against the current state of practice, with an average error less than ±0.0005. Moreover, this level of performance is extensible across the BWR fleet, and recent advancements in model architecture demonstrate remarkable resilience when new fuel types are introduced into the core.

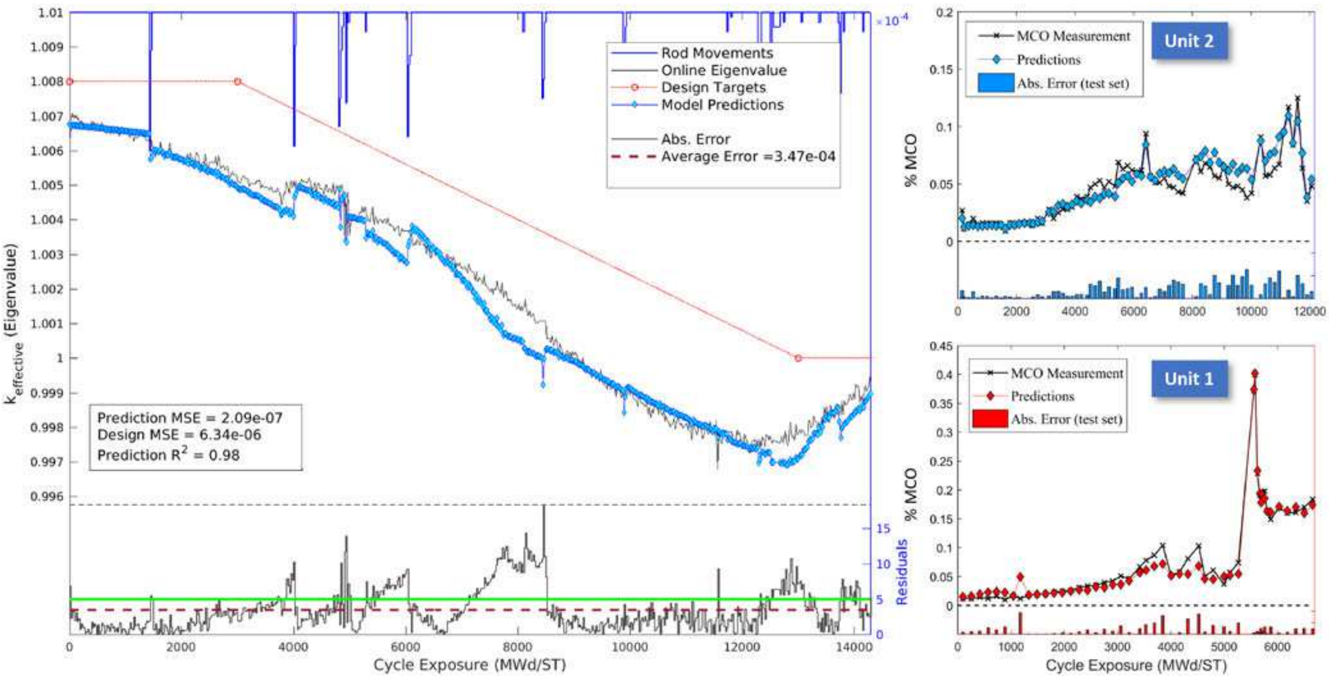


Prediction accuracy of eigenvalue and MCO ML models at two BWRs.

## *Seamless integration from reload design to cycle management*

These AI-based predictive algorithms have been turned into cloud-based platforms, MCO.ai and eigenvalue.ai, which are now fully integrated into the reload process. Each iteration of the core design, of which there are many, can be run through the platform to assess the design's impact on MCO and eigenvalue trend. Core designers can plan scenarios at will from their desks and explore hundreds of options for bundle specifications, reload batch size, loading patterns, reactivity control strategies, and so forth. Remember, these predictive models are constructed from input features derived from core conditions and core simulator outputs—all of which are projectable during the early stages of the reload design process. Consequently, these tools take those core projections and give the core designer a reliable forecast of MCO and eigenvalue behavior upward of a year before the fuel cycle even commences. In this way, the reload core design can be optimized to reduce the reload batch size and/or enrichment, lower MCO below the prescribed limits, and ensure that energy requirements are met with more reliable eigenvalue forecast.

These new capabilities extend beyond just core design. The same concepts apply to cycle management strategy evaluation. If unforeseen changes occur relative to the planned operating strategy, such as a fuel failure, unplanned

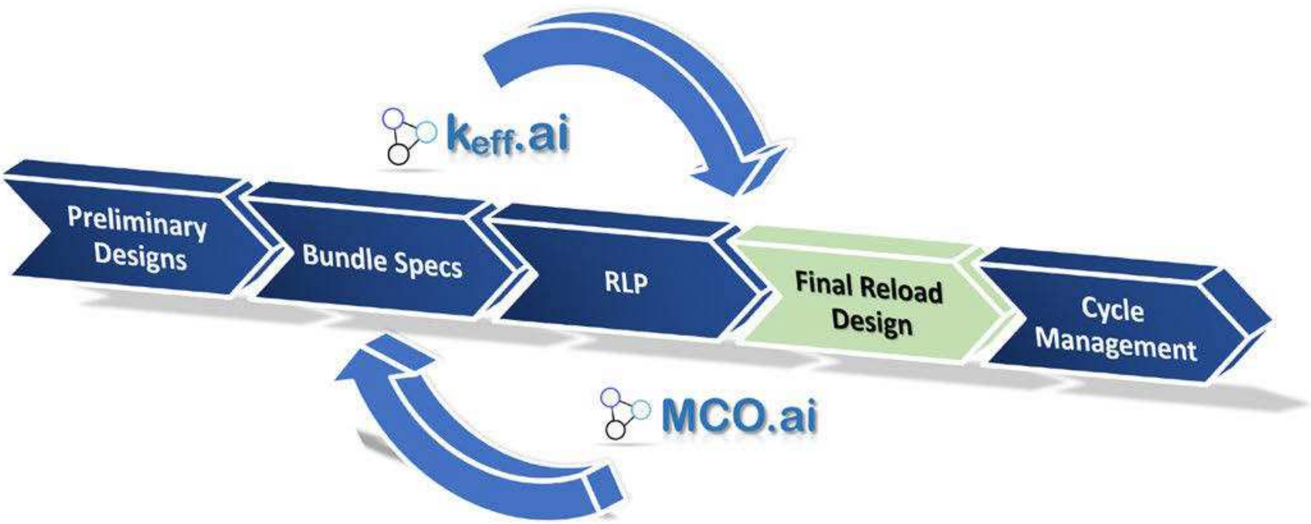


*Continued*

downtime, or startup delay, then this predictive suite can be utilized to analyze alternate operating scenarios and provide user-friendly comparisons. In fact, a fuel defect occurred at a BWR in 2019 that required power suppression control rods to be inserted through the end of cycle. The two fully inserted control rods resulted in a larger increase in MCO than normal—a step change in MCO from 0.05 percent to 0.4 percent that can be seen in the Unit 1 MCO graph on the previous page (which the model predicted with superior accuracy). This raised the question of whether a mid-cycle shutdown would be required to remove the fuel defect prior to achieving unacceptably high MCO levels. By utilizing MCO.ai, an operating strategy was devised to maintain MCO levels below the procedural limit through end of cycle, thereby avoiding a costly outage (upward of $6 million).

## The bottom line

In partnering with Blue Wave AI Labs, Constellation has achieved breakthrough levels of insight and operational predictability for two long-running problems in its BWR fleet. The use of ML integrated with the core design and cycle management processes provides fuel cost reduction, results in lower plant dose rates, protects plant assets, avoids generation revenue losses, and reduces rework, which returns hours to the business. Blue Wave AI Labs and Constellation are leading the industry with these advancements and are pursuing other applications to improve nuclear power operation and economics.

The savings of tens of millions of dollars in only a few years is just the start to uncovering the hidden treasure from the thousands of terabytes covering all plant operations across the domestic fleet. Blue Wave and Constellation are working to apply these techniques to a host of other high-value problems. These include more precise thermal limit calculations, virtual calibration and measurements, and remaining useful life of plant components that enable true condition-based maintenance strategies. Further work of this type is also proceeding in our domestic pressurized water reactor fleet. Finally, much of this insight—such as answers to the questions of sensor type and number—will be applied to next-generation plant designs. The nuclear industry is on the precipice of assuming its natural place as the central backbone of carbon-free power. AI will accelerate this ascension and deliver insights and savings at a new level. This is just the beginning.

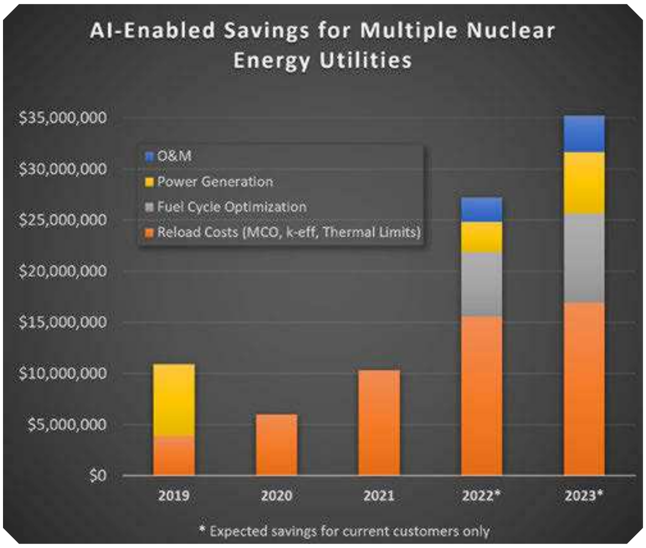


*The authors are thankful to the Nuclear Energy Institute for jointly recognizing Constellation (named Exelon Generation at the time) and Blue Wave AI Labs with a Top Innovation Practice (TIP) award for "Moisture Carryover (MCO) Predictions through Neural Networks" at the recent NEI 2021 Annual Meeting. The prestigious award in the nuclear fuel category recognizes creative ideas that have substantial impact on improving the safety and reliability of nuclear energy. The concept of MCO.ai was also selected to be part of the Electric Power Research Institute's Plant Modernization Toolbox as a Modernization Technology Assessment. EPRI's Plant Modernization Toolbox is a resource to facilitate decision making and execution of the modernization process at nuclear power plants. It includes a variety of tools and aids to assist nuclear plants to identify and evaluate cost savings from technology and process improvements.*

*J. Thomas Gruenwald is the chief operations officer and cofounder of Blue Wave AI Labs.*

*Jonathan Nistor (jonathan@bwailabs.com) is the chief science officer at Blue Wave AI Labs.*

*James Tusar is the senior manager of BWR Design and Methods in Constellation's Nuclear Fuels Department.*