eHANDBOOK

CONTROL

PROMOTING EXCELLENCE IN PROCESS AUTOMATION CONTROLGLOBAL.COM

AI in Industrial Operations: The Path to Augmented Efficiency

Perhaps no topic has captured the interest of process control engineers than the growing opportunities, and challenges, of artificial intelligence (AI) in industrial operations. While we're still a long way from surrendering control of operations to AI, the chance to augment operations and increase efficiency are too great to ignore. While the world catches on, Control editors and contributors have been following the rise of AI in industry for many years.

From features to news to podcasts, here's a sampling of our coverage with plenty more to come.

TABLE OF CONTENTS

Artificial intelligence expected to add $10.5 billion to manufacturing by 2033

Generative AI is growing because participants are building potential use cases and scaling from creating new designs to overhauling production processes

By Jim Montague

Investments in generative artificial intelligence (AI) are expected to increase manufacturing revenues by $4.4 billion during 2026-29, and continue onward to reach $10.5 billion by 2033, according to a recent report, “Generative AI use cases in manufacturing,” by ABI Research. It reported Aug. 30 that generative AI is growing because participants are building potential use cases, and scaling from creating new designs to overhauling production processes.

“Generative AI has growth that will derive from functionality and use cases across market verticals,” says James Iversen, manufacturing and industrial industry analyst at ABI. “Deployment of generative AI will come in three waves as the technology matures, with manufacturing seeing the largest revenue growth during the second and third waves. During the second and third waves of adoption, generative AI will be deployed into four domains of manufacturing—design, engineering, production and operations.”

ABI projects the design field will see the fastest mainstream AI deployment. It adds that use cases, such as generative design, manufacturing bill of materials (MBOM) and electrical bill of materials (EBOM) reductions, have already generated design products by companies, such as Siemens and Microsoft. ABI adds that engineering, production, and operations use cases will take longer and require further maturity from generative AI providers due to the task complexity and required model training.

Use cases for generative AI in manufacturing can be compared by looking at expected time to value (TTV) and return on investment (ROI). The top performers for the four domains are:

- **Design**—generative design and part consolidation;
- **Engineering**—tool path optimization and part nesting;
- **Production**—root-cause analysis of product quality and correcting buggy software code; and
- **Operations**—inventory stock and purchasing period management, as well as employee work path optimization.

API adds that manufacturers and manufacturing software providers initiating use cases include BMW, Boeing, ByteLake, General Motors, Markforged, Nike, Nvidia and SprutCam X. They're aided by generative AI companies, such as Nike's Celect, Gradio, OpenAI, Retrocausal, Work Metrics and Zapata AI.

API concludes that manufacturers and manufacturing software providers should prioritize top-performing use cases because they yield the highest returns and can be easily built out with existing generative AI capabilities. "Starting from the ground up, implementing these use cases will lay the groundwork for more extensive use cases," explains Iversen. "It's important not to jump the gun and develop high-functioning use cases that will see little implementation because trust in generative AI will need to be built up before overhauling significant portions of current manufacturing operations."

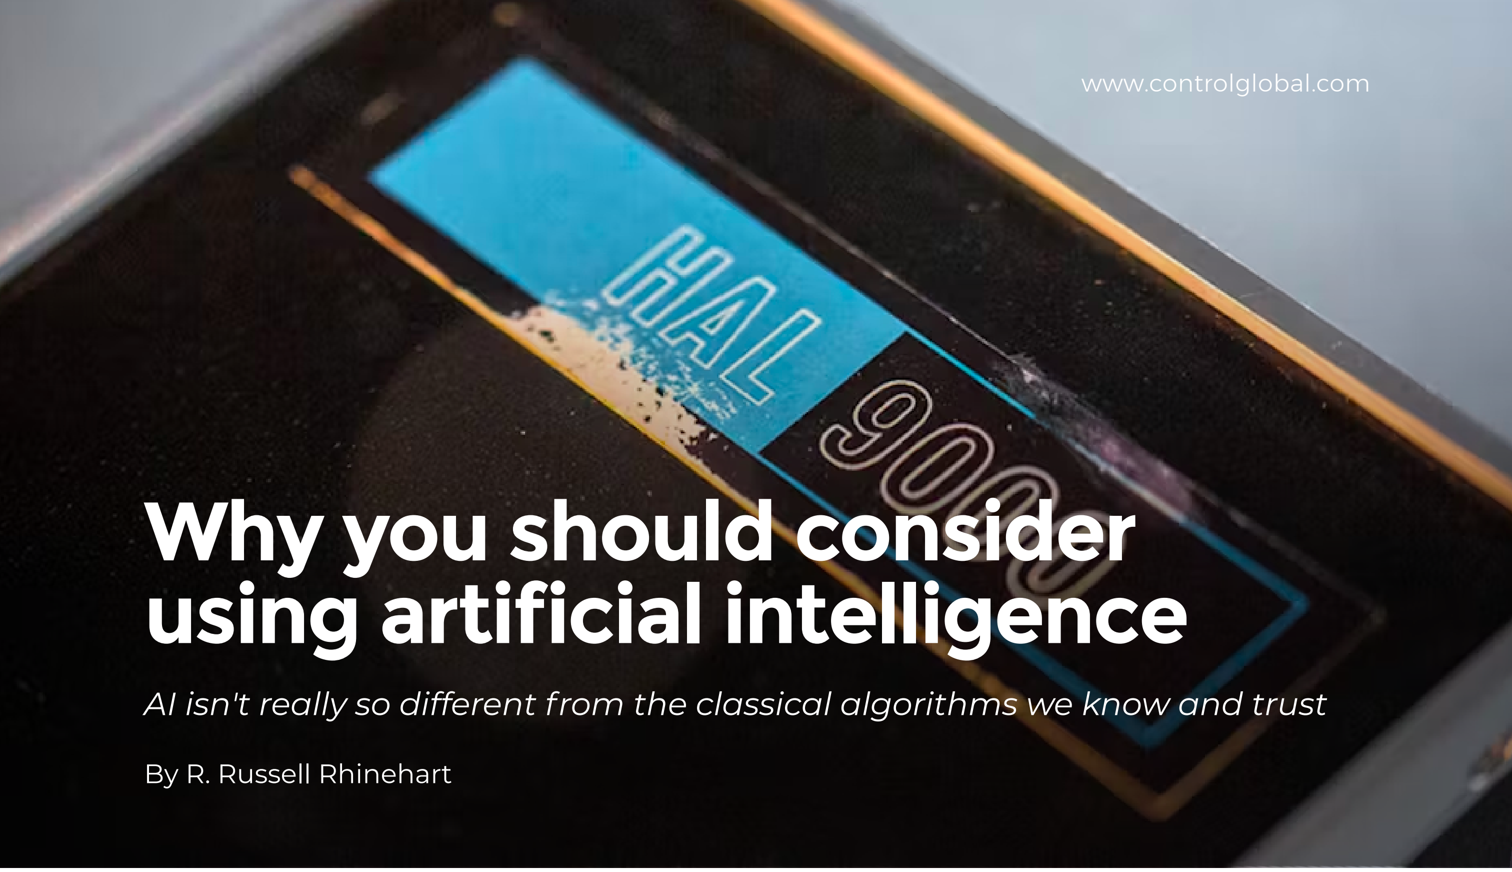

Why you should consider using artificial intelligence

AI isn't really so different from the classical algorithms we know and trust

By R. Russell Rhinehart

We're comfortable using regression models based on a power series such as $y = a + bx + cx^2 + \ldots$ Here, the model structure is the sum of power-series terms and the a, b, c coefficients are linear in the model (given an x-value, the calculated y-value is linearly proportional to the coefficient value). We get the model coefficient values by best matching the model to the data; and with this structure, linear regression is easy. Simple algebra is used to calculate the steady-state y from steady x conditions. If a quadratic model doesn't seem to have a good fit to the data, try a cubic model. We admit the model isn't the truth about the x-to-y relation, and that it takes human judgment to balance goodness-of-fit against model complexity. Data could be generated designing experiment protocols, or it could be acquired from historical trends.

We also accept first-order-plus-dead-time (FOPDT) models of a transient response, which are often presented in Laplace transform notation such as:

$$\hat{y} = \frac{Ke^{-\theta s}}{\tau s + 1}\hat{u}$$

Here, the model structure represents a first-order lag after a delay, and the modeled y-value isn't linearly dependent on the coefficients of τ or θ. We don't solve for the y-value in that Laplace notation, but use numerical methods to solve the differential equation it represents. Again, we adjust model coefficient values to best match model to the data; but here, we need to use nonlinear regression. If you want a better fit to the data, perhaps use a second-order model (SOPDT), or schedule the model gain, time-constant or delay. Data might be generated by the reaction curve method—a single step-and-hold in the measured variable (MV) and observation of the transient controlled variable (CV) response.

Finite impulse response (FIR) models of transient behavior are common in model-based control. Use a pseudo-random binary sequence to make sequential changes in the MV, observe the CV, and adjust impulse heights in the model to best fit. Use human judgement to specify the vector length, sampling interval and model inputs that provide the desired balance of fit to complexity. Linear algebra techniques are used to solve the matrix-vector models.

Those three modeling approaches are unique; each has its own mathematical form, solution procedure, protocol for generating data, regression algorithm, and criteria for goodness. A user wouldn't use a steady-state model to try to match a transient response. A user can't use methods for one to make another generate a right answer. If the match is not good enough, the user chooses what additional complexity to include in the model.

AI UNDERUTILIZED

There are many other modeling approaches in the category of artificial intelligence (AI). They use modeling forms that are realizable with modern computing tools. AI techniques offer improved flexibility as modeling tools, but I believe they're underutilized.

My experience is that industrial managers have three objections to any form of AI. First, there was a view that came from those who attempted to implement neural networks (NN) or fuzzy logic (FL) as young engineers, after being snookered by the promotional hyperbole of those selling the new tools. However, because techniques and methods were unfamiliar, or the promise was exaggerated, they couldn't get it to work. And, rather than admit "I was inadequate to get it to work," they claimed "It doesn't work." They kept this view as they eventually got promoted to managerial positions, and many years later blocked FL and NN applications based on prior notions.

Second, there's the adverse reaction to names from the research/science community that created AI approaches. Fuzzy logic sounds like faulty thinking. Artificial intelligence sounds like an independent sentient presence with a possibly adverse agenda that can't be managed by humans (Consider HAL from the movie 2001: A Space Odyssey.) And a neural network sounds like a mysterious brain-like thinking method of calculation. Instead of the term regression to mean fitting model to data, they use anthropomorphic terms such as learning, training or memorization, making AI seem to be a form of intelligence. By contrast, FL is just a way to implement human logic, and it's not faulty. In spite of the hyperbole of researchers and movie producers, there's no sentient presence behind an AI computation. Neural networks are not unknowable. The code and computation is as simple as any regression modeling that we use in multivariable statistics or model-predictive control.

The third manager objection I've encountered is a fear of the unfamiliar. How can one be confident that the models (understanding) are complete and comprehensive—and reliable? How can one be sure that AI is properly adapting to a new situation and understands its context? Managers are in 'responsible charge,' and have the same fear about strange techniques as they do about how new engineers assess situations and make recommendations. Interestingly, the same questions can be cast at any modeling approach, but practice has made us comfortable with classic modeling approaches. I now think we have substantial application evidence that AI is as safe and legitimate as any modeling approach.

So develop your skills and consider NNs when traditional regression models don't give desired results. Consider FL to automate human decision-making. First, understand the techniques and associated methods to develop the models. Use procedures designed for the technique, not for an alternate modeling approach.

As a message to vendors, I think it's important to not let excitement about new techniques lead to statements that make practitioners think it is either a cure-all or ready for an untrained person to apply. Vendors need to provide details of the methods, and complete instructions to users about what to do and not do. I think vendors shouldn't use the anthropomorphic terminology of the researchers. They should return to the familiar terminology of the users.

AI likely to augment humans, not replace them

Intel's Rita Wouhaybi together with co-panelists from Microsoft, AspenTech and Emerson, discussed the rapidly evolving capabilities of artificial intelligence (AI) and the role it will play in plant automation.

By Len Vermillion

Automation intelligence (AI) and machine learning (ML) certainly aren't new concepts to process control engineers. These technologies have been in play in industrial processes for 30 years or more, but they've largely been the realm of the technologist or data scientist. Today, AI has blown the minds of many people, and large language models (LLMs) have exposed generative AI to just about anyone with a Wi-Fi connection and an internet browser. Let's just say it has more than a few people thinking about the possibilities and the challenges, as well as the dangers of AI.

"We have been working on this for many years and believe it will be a big transformation for all of us," said Jose Valls, CTO of manufacturing for Microsoft, during a thought-provoking panel discussion called, "Forging Tomorrow: Unleashing AI Benefits" during Emerson Exchange Immerse this week in Anaheim, Calif. "It's here and it's not going to stop," Valls said.

Valls is far from the only one working on generative AI to take notice. "In my tenure, I don't think I've ever seen something that has so much potential," added fellow panelist Rick Kephart, vice president of technology at Emerson's Power and Water Solutions business, who has worked in industrial control for more than 30 years. "I've become a believer, and I wouldn't have said that last year."

AI is winning over many people in the industrial sector, just as it is in a plethora of other industries. At the same time, it stokes fears ranging from the practical—jobs losses, mistakes in production—to the fantastical—a dystopian society run by computers. (Perhaps thanks to decades of Hollywood movies that tend to turn toward the dark side for dramatic effect.)

No need to worry about that just yet, the panelists pointed out. We're a long way away from a total surrender to autonomous manufacturing and processing plants. Rather, the focus these days remains on the opportunity to use AI to augment human productivity and efficiency. It can also help increase cybersecurity and super-charge real-time diagnostics and predictive maintenance in the plant.

AUGMENTATION OR AUTONOMY?

"I'm really interested in the runtime and diagnostics," Kephart continued, pointing out two areas he's seen significant benefits already: machinery health and monitoring and reacting to off-normal conditions, much like a digital twin. "I think we are a long way away from taking the human out of [the plant]," he added. "We'll be needed even if to moderate between competing control functions. I'm a little skeptical about the autonomous plant."

Rita Wouhaybi, senior AI principal engineer, Intel, also isn't concerned about losing the human touch at this point. "I love the fact that we're talking about AI assisting humans," she said. "Playing the role of assistant is the right attitude."

For AI to become humans' helper, it must speak the language of humans, she added. That's where LLMs entered the party and turned it into a big bash. "It has to be present where anyone can consume it," Wouhaybi continued. "At the end of the day, is it deployed in my pocket on my phone? Where is the value being delivered?"

Like Kephart, she sees value in diagnostics right now. She added that predictive opportunities are exciting for the near future, saying LLMs are good at predicting the next move. "They are basically saying what the human would do next," she said.

Nithiya Parmeswaran, vice president of product management at AspenTech, also said the value of AI must be established for industrial processes before companies jump in the water with both feet. "In the industrial space, we need to talk more about the value and less about the technology," he said. "We need to explain to the end user so they can understand the cost and why they should be using it."

While the promise is off the charts, the hindrances are still many, with data at the top of the list. The dirty truth about AI is that it can only be as good as the data it is trained on.

The panelist said LLMs must be trained on specific plants, much like digital twins, to be successful. The idea is for the AI to know the plant just as a human worker would, so that it can function as a “shirt tug” to prod humans to react. There’s no doubt the holy grail would be for AI to handle issues itself, but we’re not there yet, they agreed.

Another hindrance, they said, is that AI will have to auto-learn to hit such heights. At the moment, it is reliant upon data scientists to have clean data, but as anyone working with data knows, perfect data is non-existent, so until AI can learn to sort good data from the bad, humans will remain in charge of the plant.

At the end of the day, AI is ready to be a super-worker, but don’t expect it to be the boss anytime soon.

Leveraging AI for enhanced asset performance management

Artificial intelligence and other intelligent device management tools can better optimize industrial processes

By Ian Verhappen

Asset performance occurs at many levels—from macro-level optimization of facilities to individual processes and pieces of equipment, including the individual sensors and actuators used to control the process. Many automation professionals associate asset management with the lowest level—keeping track of individual field device quantities, configuration, calibrations and related work orders.

Most field devices installed in today's facilities include some form of digital communication, and can communicate more information than the process variable to support this lowest level of asset management. Communication devices use some form of onboard diagnostics, such as heuristic rules and pattern recognition, to determine status. Device status, which is analogous to a common trouble alarm for the device, is one of the first information pieces to consider sharing. If device status is communicated to support NE 107, it classifies the status as one of four broad categories that indicate whether the device is still usable or out of service. This indication immediately informs the operator about the urgency for action and potential impact to closed-loop control.

ISA-108 is working with IEC SC65E WG10 to develop Intelligent Device Management standards on procedures and best practices to allow intelligent devices to improve asset utilization and effectiveness at this lowest level, which is the foundation for other asset operations including basic regulatory control.

Despite being fundamental to the control process itself, the search for much larger returns focuses on large equipment, particularly rotating equipment such as pumps and compressors. These devices consume significant power, have many parts that can fail, and can cause unscheduled outages that impact the processes to which they're connected.

Artificial intelligence (AI) and associated pattern recognition tools can identify the key inputs used by lower-level diagnostic models, and then work in real-time to improve performance management by providing predictive, diagnostic and forecasting analytics. They work together to detect emerging issues, identify causes, and indicate required timeframes for action. These models require maintenance to accurately account for changes in equipment, including new process conditions such as start-up, shutdown, start-of-run versus end-of run, and changes in the process (feed, blends, etc.)

If the objective is to optimize performance of larger, facility-level assets, external factors impacting performance such as cost of power and customer demand must be incorporated into the model.

Asset performance models are developed using a wide range of tools and inputs. Understanding the relevance of the model, boundary conditions and principles of data science help identify relevant and irrelevant data that are keys to success. These include digital twins, which when fully implemented, can use relevant process data such as failures and analysis of root causes. They can be used as inputs to improve and refine the model.

Most larger asset performance models can be complex, so they're run "offline" and periodically, such as once per hour, once per day or even only on-demand.

Other asset performance models are meant to be used more like optimization routines for a facility or process, which means they'll run much more quickly and with fewer computing resources. If used this way, the models can be simple enough to run in real-time to provide guidance to operators and help optimize the process.

Multiple suppliers of predictive asset management services claim performance improvements of 2% to 6% increases in availability, 10% to 40% reduction in reactive maintenance, and a 5% to 10% reduction in inventory costs. Similar improvements of up to 30% can be achieved by applying these tools in an asset utilization/capacity.

Asset performance management is about more than field devices. Though the information provided by sensors and actuator diagnostic tools are important inputs to the ability to control the process, applying tools such as AI and resultant models to make intelligent decisions about the process at all levels has proven benefits.

Transparency, supervision still needed for mission-critical AI

Pete Diffley
Senior Manager
Global Partnerships
Trihedral Engineering

With artificial intelligence (AI), machine learning, digital twins and soft-robots inserting themselves into most of the systems we interact with today, it's important to realize that, while logic-enabled, pattern-recognizing algorithms demonstrate some aspects of intelligence, they're not yet a panacea for mission-critical decision-making.

According to Pete Diffley, long-time veteran of process manufacturing's front lines and current global partnerships lead for VTScada by Trihedral, such algorithms are the right tools for many jobs, and can bring new, automated efficiency to formerly manual tasks. But in some situations, unforeseen biases can skew results in a number of unanticipated directions. Control sat down with Diffley to discuss some of the sources of these biases, and what steps can be taken to ensure that our mission-critical systems don't lose their way.

Q: Are the terms AI and machine learning being overused?

A: Especially in a broader, non-industrial context, if any tool or application includes logic in it these days, someone, somewhere is trying to say, “That’s AI or that’s machine learning,” and usually it isn’t either. Often, it’s a sensor and an output related by a simple algorithm. It’s automation. AI and machine learning are real phenomena, of course, and aspects of intelligence are creeping into all kinds of systems, but real intelligence is much bigger than an algorithm with a bit of logic attached to it. Algorithms that appear to “learn,” which can adjust their response based on the data sets they’ve seen—their “experience” of the world—are I think one of the reasons that the term AI is so overused. We identify learning with entities that have intelligence, be they people or guinea pigs.

Q: Bias is another very anthropomorphic term that affects not only people but also the AI or machine-learning algorithms they create. Can you talk a bit about the sources of bias that can assert themselves?

A: Bias can be introduced at the very start, when it comes to selecting the best algorithm for a particular application from the proliferating number of commercially available and open-source options available. Familiarity or inclusion of a particular algorithm in one’s toolkit may well exclude other, more suitable alternatives. Algorithm training presents another opportunity to introduce bias. If the data sets presented to the algorithm don't represent the full range of expected process conditions, the algorithm may find itself on uncertain ground. Imagine training an algorithm to recognize dogs, but only using Chihuahuas. Will it mistake the first Great Dane it sees for a horse?

AI or machine-learning systems are also subject to the same sort of selective blindness that affects humans at times. For example, if I present a data set to person consisting of the numbers 2, 4, 6, 8 and 10, a logical person will likely jump to the conclusion that the data set is incrementing by two. But what if it’s just the set of positive integers that 5 and 7 also belong to? ? As humans we tend to seek out patterns, so we can process things more quickly. An algorithm—like a person—may errantly identify a pattern that then becomes hard to “unsee” in the data.

Q: What steps can be taken to minimize the possibility of AI or machine-learning systems acting in unpredictable ways?

A: A number of high-profile incidents have demonstrated just how important it is that the decision-making criteria of these systems be well-understood by their human operators. It’s obviously not a huge issue if Amazon gets its purchase recommendations wrong, but when we're talking about mission-critical systems—say, a highspeed turbine or an offshore oil rig—the stakes are much higher.

The crashes that resulted from problems with the Boeing 747 Max 8 flight control system, for example, were at least partly to blame on a poor understanding of how the system worked by the pilots. The pilots were confused and didn't understand what steps the system was taking or why. The lesson to be learned from this is that you can't operate mission-critical AI as a black box—at least not yet. Human operators need to have full transparency into the criteria the algorithm is using to make a particular decision or recommendation.

In another aircraft example with a happier ending, the triple redundant pitot tubes for airspeed measurement on an Airbus model had to be covered while on the ground to keep the local bee population from making themselves at home in them. Somehow, multiple work processes and inspections by numerous critical personnel failed to ensure that the caps were removed before take-off—confirming a human tendency toward selective blindness. And as the plane accelerated down the runway with caps still in place, the plane's systems alerted them that the plane had no airspeed measurement. In this case, the plane's systems, presuming a fault in the airspeed sensors, guided the pilots to its GPS readings instead, also using stresses on the wings to estimate the plane's speed. Turns out, the plane didn't need airspeed indication via the pitot tubes—it had other redundant systems that could fill in the gaps. Fortunately, they successfully landed the plane and marked an incident, not a tragedy.

These two examples show that nothing is totally foolproof and that even experienced personnel can become blind to things that should be obvious in certain situations. A friend of mine that has an amazing solution in the cybersecurity sector rightly reminded me that it's all about framing when it comes to algorithms and AI. Knowing when it's appropriate to use it, the maturity level to which it's adopted, and the type of algorithms being used. It's also essential to not allow “incremental creep” to set in over time. Something can be deviating minutely over days, weeks, months or even years, and wind up a long way from where intended. Having a frame of reference of where you are versus where you thought you should be is extremely important, especially with AI.

Q: What lessons can process manufacturers and other operators of mission-critical infrastructure learn from these examples?

A: I'd say the biggest takeaway it's not leaving AI or machine-learning algorithms to operate as black boxes with everything hidden inside. You can keep saying “train the system,” but there’s still unavoidable bias. Have the system share the logic it used to reach a particular decision. That's transparency, and of course it’s complicated to deliver. That's why you don't get it often.

AI decisions themselves aren’t necessarily any better—or worse—than those made by humans. People often make decisions based on the wrong way of looking at data because everything is open to interpretation. When operations are mission critical, both systems and people need to be able to say, “I'm basing my decision on these factors.” Then someone should be in a position to challenge it, and maybe AI could challenge that in turn. So, perhaps you have several layers of AI. One layer disputes the conclusion of another layer, followed by yet another layer of artificial arbitration—which now and then is presented to a human, who checks in and says, “Okay. We haven't quite reached the Skynet singularity."

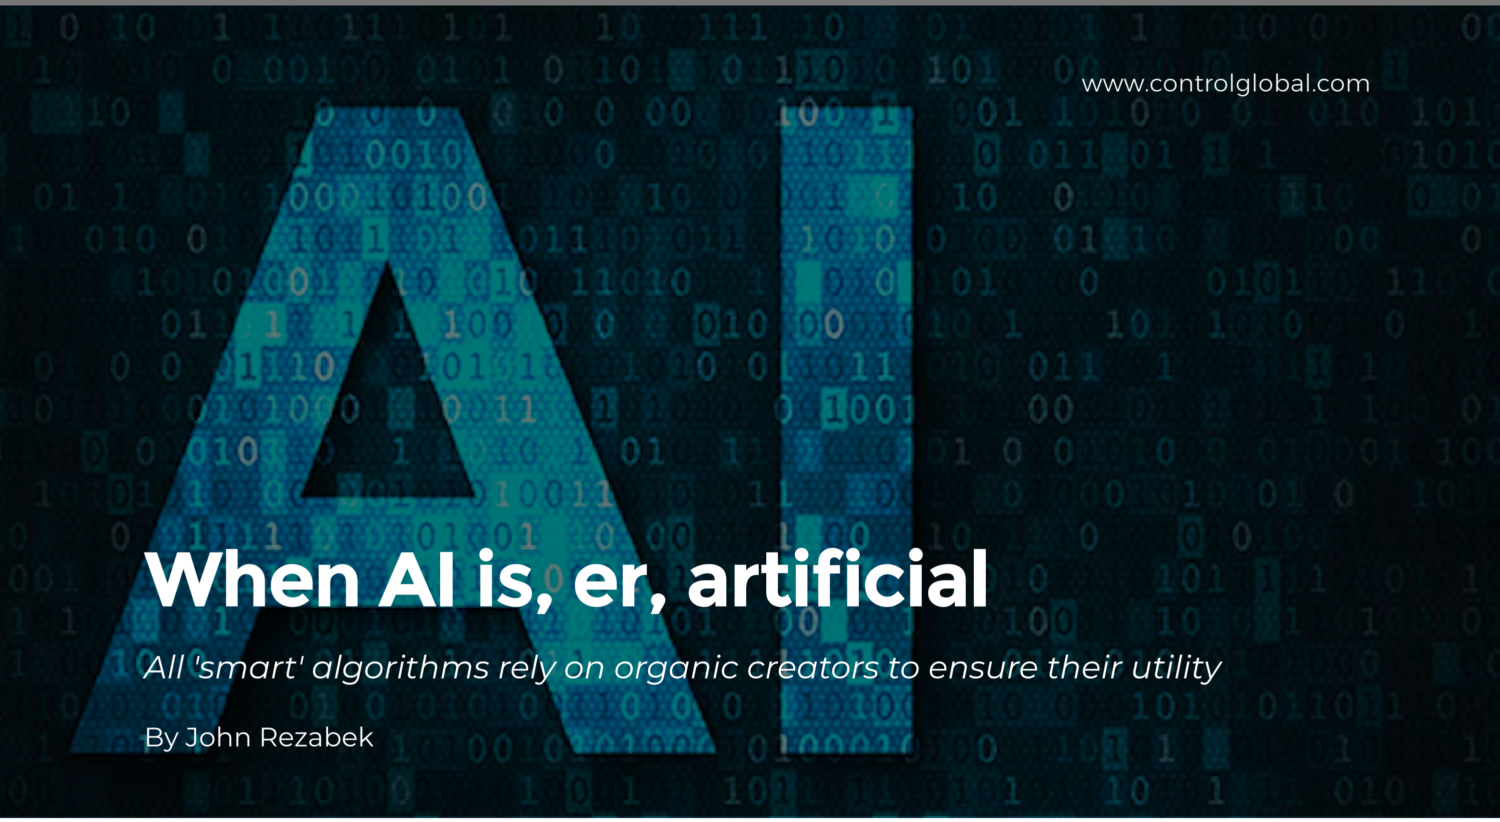

When AI is, er, artificial

All 'smart' algorithms rely on organic creators to ensure their utility

By John Rezabek

Intelligence that's "anything but artificial" boast the ads for luxury, battery-powered cars. AI engines may be a boon to consumer marketers, who are getting better at mapping the robotic habits and appetites of their organic end users. But should we be tempering our expectations at all as we contemplate how such algorithms are applied to industrial processes?

A decade ago, our chief console operator, Chaz, witnessed the newly deployed model predictive control (MPC) application making moves in anticipation of changing measured variables and measured disturbances. He liked to call it "Spock's Brain," which fans of Star Trek TOS ("the original series," for the uninitiated) will recall featured a fetching alien intruder, who removes the science officer's brain to replace a worn-out brain controlling a subterranean HVAC. The newly lobotomized Spock accompanies the usual crew, who eventually discover his new role, ensuring his captors could remain cozy and comfortable in their 1960s go-go boots.

The creator and writers of science fiction have imagined AI like the conversant ship's computer on the Enterprise and Knight Rider's KITT for decades. We're experiencing it firsthand now as we interact with Siri, Hey Google and Alexa, seeking navigation guidance, movie trivia, sports updates and even tuning suggestions for pH loops.

Just as the MPC controls were just doing math, there's no personality churning through the algorithms to answer questions on behalf of the Google or IOS user. These services are astonishing in their abilities to index and retrieve seemingly the sum of all human knowledge in seconds. However, they're balanced by the times they get it entirely wrong—no, I don't want to call the mayor of Tunis, I want to know if Meijer supermarket has tuna.

So, when a Bayesian optimization engine recently prompted the operators to drive the unit into perilous territory, Chaz likened it to another Star Trek TOS episode, "The Ultimate Computer," in which the brilliant scientist Dr. Richard Daystrom's M-5 computer was determined to wipe out all the humans in pursuit of self-preservation. The moral was that even the most thoughtfully crafted logic could lead to unintended consequences. It's incumbent on AI creators to imagine every if-then-else that might befall their programming. The engineers of self-driving cars, for example, must consider the torments their machine will encounter from humans determined to befuddle it. Star Trek enthusiasts will recall how Captain Kirk foiled the M-5 with its own logic. How prescient were the sci-fi writers and producers of the mid-20th century?

TO BOLDLY GO

Data-lake mariners are also boldly going where no one has before—and have justifiable enthusiasm for the prospect of finding a new world of golden treasures in the murky depths. But, like the artificial lake made by the U.S. Army Corps of Engineers in the next county, no one bothered to clear all the stumps. Even though the lake may be vast and polluted by trash, one can hypothesize with some confidence that there are valuable insights to be found. Many providers promise their algorithms can dredge through the deposits of sediment, revealing correlations and forecasting trajectories.

Whatever you call your algorithm, George Box's maxim still applies: "All models are wrong, but some are useful." They exist almost exclusively in cyberspace, while their real-world counterparts are often haunted by all manner of curious, unmeasured disturbances and dynamics. Twinning something as simple as a waste heat incinerator and boiler can still run across relationships that defy "useful" modeling. It's not uncommon for carbon monoxide in the stack to react faster than the temperature sensors. Just try correlating that.

The process industries are no strangers to predictive controls and optimizers. Contemporary AI engines and methods have new promise as we channel more pervasive sensing and other flotsam into the data lake. But let's keep in mind that robots excel at being robotic, and fantasies like Knight Rider's KITT or the Enterprise's guidance computer require organic intelligence to ensure their utility.

Exploring the impact of generative AI on industrial operations

A Control Amplified podcast with Dustin Johnson, CTO of Seeq

By Control Amplified

Generative AI is a topic on the minds of a lot of people these days. The emergence of automation intelligence, especially over the last two years, has disrupted the way industry thinks about future operations. Although scary to many people, generative AI ushers in new and exciting possibilities for manufacturers and process industries. It has the potential to reshape the way they analyze data, optimize operations and make critical decisions.

To learn more, Control editor-in-chief Len Vermillion hosts Dustin Johnson, CTO of Seeq, which provides analytics software to process industries.

Listen to full podcast **HERE**