

Technology Scan

Focus: Artificial intelligence for energy transition

ASIA-PACIFIC

AUSTRALIA

AI to make green ammonia greener

To find the best catalyst for green ammonia, researchers were staring down 8000 lab experiments. With AI, they only needed 28. Scientists and engineers at the University of New South Wales (UNSW) Sydney, who previously developed a method for making green ammonia, have now turned to artificial intelligence and machine learning to make the process even more efficient.

The UNSW team discovered a way to make ammonia from air and water using renewable energy, at about the same temperature as a warm summer's day. Dr Ali Jalili, with UNSW's School of Chemistry, says while the original proof-of-concept demonstrated that ammonia could be created entirely from renewable energy, at low temperatures and without emitting carbon, there was still room for improvement. For example, could it be produced more efficiently, using lower energy, less wasted energy, and producing more ammonia?

To answer these questions, the team needed to find the right catalyst – a substance that speeds up the chemical reaction without being consumed by it. But the best catalyst would need a combination of these metals, and if you do the maths, that turns out to be more than 8000 different combinations. The researchers fed a machine learning system information about how each metal behaves and trained it to spot the best combinations. That way, instead of having to run more than 8000 experiments in the lab, they only had to run 28. The winning combo was a mix of iron, bismuth, nickel, tin, and zinc. While the researchers were expecting some improvement in the process of producing green ammonia, this new five-metal catalyst exceeded even their most optimistic expectations.

Known as Faradaic efficiency, high efficiency scores mean the process is more sustainable, cost-effective, and scalable, which is crucial for making green ammonia a viable alternative to fossil-fuel-based methods. Dr Jalili says his team was able to make ammonia this way at an ambient 25°C, less than 10% of the temperature required to make ammonia the conventional way via the Haber-Bosch method.

This low-temperature, high-efficiency approach makes green ammonia production viable and scalable. We believe it can compete directly with electrified Haber-Bosch and even fossil-based routes, creating a realistic pathway for truly green ammonia. The goal is that one day soon, farmers will be able to produce ammonia for fertilisers onsite, at low cost and low energy, eliminating the need for delivery via transport routes, further reducing the carbon footprint of ammonia production.

<https://www.unsw.edu.au/>

CHINA

AI improves fusion reactor safety and performance

A research team led by Prof. Sun Youwen at the Hefei Institutes of Physical Science of the Chinese Academy of Sciences has unveiled two artificial intelligence systems designed to enhance the stability and efficiency of fusion experiments. Their results appear in the journals *Nuclear Fusion and Plasma Physics and Controlled Fusion*.

Fusion energy promises clean, virtually inexhaustible power, but future reactors must operate reliably to prevent damaging disruptions and maintain precise plasma confinement. Disruptions are sudden, intense events that threaten reactor integrity, while maintaining high-performance confinement states is critical for sustained operation.

To address these issues, the team built two specialized AI platforms. The first, a disruption predictor, employs decision tree models to identify early warning signals of disruptions triggered by locked modes, a common plasma instability. Unlike opaque black-box algorithms, this model is interpretable, pointing to the physical causes behind its predictions. In trials, it successfully issued warnings 94 percent of the time, with alerts arriving 137 milliseconds before the disruption-leaving enough time for countermeasures.

The second AI system focuses on real-time plasma monitoring. Instead of relying on separate models to classify confinement states such as L-mode and H-mode and to detect edge-localized modes (ELMs), the researchers developed a multi-task learning framework that performs both simultaneously. This approach increased both accuracy and resilience, achieving a 96.7 percent success rate in recognizing plasma conditions.

Together, these innovations advance the prospects of next-generation fusion reactors by boosting safety, improving performance, and contributing to deeper insights into plasma dynamics.

<https://www.spacedaily.com/>

INDIA

AI-powered solar manufacturing line

Pralhad Joshi, Union Minister of New and Renewable Energy, launched India's first AI-powered solar manufacturing line at Goldi Solar's newly built, state-of-the-art facility in Kosamba, Surat, Gujarat. This landmark achievement marks a major leap in India's renewable energy sector, integrating cutting-edge AI-driven processes to revolutionize solar PV module production. The AI-powered facility enhances precision, scalability, and efficiency, redefining solar manufacturing and strengthening India's transition toward a net zero future.

The AI-powered facility, with an impressive 14 GW planned production capacity, integrates several industry-first innovations designed to redefine solar manufacturing. These include high-speed stringers leveraging AI-driven automation to achieve a remarkable production capacity of 10,000 cells per hour, ensuring consistency and precision, minimizing errors, and reducing material wastage. The facility also features AI-powered AOI (Automated Optical Inspection) systems for real-time quality control, using computer vision to detect and eliminate defects at a microscopic level.

The facility further incorporates a fully integrated pre-lamination section with an auto EVA laying robot, which optimizes the placement of encapsulation layers with minimal human intervention, reducing inconsistencies compared to traditional manual processes. Additionally, it includes AI-powered 3-Display EL and visual (front and back) inspection systems, which ensure that only the highest quality solar modules reach the market, enhancing reliability and durability.

Advanced automation is further demonstrated through a 3-stage auto laminator with double-side heating (electric and oil-based) that enhances panel strength and longevity, an auto ICB elevate-lifting system with Teflon removal to optimize efficiency, and an inline double-layer HI POT tester to ensure superior insulation and electrical safety. Additionally, the facility incorporates an intelligent 8-bin auto sorting system for precise module classification, leveraging AI for real-time adjustments, and a state-of-the-art MES system for live module monitoring and tracking, ensuring complete transparency and operational efficiency.

Capt. Ishver Dholakiya, Founder and Managing Director of Goldi Solar, commented on the milestone, saying, "Goldi Solar is proud to be the first solar company in India to integrate Artificial Intelligence into PV module manufacturing. Our planned 14 GW advanced facility is a testament to our vision of a net-zero India and reflects our unwavering commitment to innovation, sustainability, and excellence."

<https://constructiontimes.co.in/>

JAPAN

ML to enhance nickel catalysts for CO₂-to-methane conversion

Researchers at Tohoku University have utilized explainable machine learning to uncover critical factors that enhance the performance of nickel-based catalysts in the conversion of carbon dioxide into methane. The study highlights how data-driven approaches can inform catalyst design, contributing to advancements in carbon recycling and sustainable energy solutions.

The team applied machine learning techniques to analyze the properties of nickel catalysts used in CO₂ methanation, a process that transforms carbon dioxide into methane—a potential renewable energy source. By leveraging explainable AI, researchers identified specific characteristics and conditions that optimize the catalytic activity. This approach not only provides insights into improving existing catalysts but also offers a framework for developing new materials tailored for efficient carbon recycling applications. The findings underscore the role of artificial intelligence in addressing environmental challenges through innovative material design.

<https://www.geneonline.com/>

THE REPUBLIC OF KOREA

Energy-efficient NPU technology

Researchers at the Korea Advanced Institute of Science and Technology (KAIST) have developed energy-efficient neural processing units (NPU) technology that demonstrates substantial performance improvements in laboratory testing. Their specialised AI chip ran AI models 60% faster while using 44% less electricity than the graphics cards currently powering most AI systems, based on results from controlled experiments.

To put it simply, the research, led by Professor Jongse Park from KAIST's School of Computing in collaboration with HyperAccel Inc., addresses one of the most pressing challenges in modern AI infrastructure: the enormous energy

and hardware requirements of large-scale generative AI models.

Current systems such as OpenAI's ChatGPT-4 and Google's Gemini 2.5 demand not only high memory bandwidth but also substantial memory capacity, driving companies like Microsoft and Google to purchase hundreds of thousands of NVIDIA GPUs.

The core innovation lies in the team's approach to solving memory bottleneck issues that plague existing AI infrastructure. Their energy-efficient NPU technology focuses on "lightweight" the inference process while minimising accuracy loss—a critical balance that has proven challenging for previous solutions. The technology centres on KV cache quantisation, which the researchers identify as accounting for most memory usage in generative AI systems. By optimising this component, the team enables the same level of AI infrastructure performance using fewer NPU devices compared to traditional GPU-based systems.

The KAIST team's energy-efficient NPU technology employs a three-pronged quantisation algorithm: threshold-based online-offline hybrid quantisation, group-shift quantisation, and fused dense-and-sparse encoding. This approach allows the system to integrate with existing memory interfaces without requiring changes to operational logic in current NPU architectures.

The hardware architecture incorporates page-level memory management techniques for efficient utilisation of limited memory bandwidth and capacity. Additionally, the team introduced new encoding techniques specifically optimised for quantised KV cache, addressing the unique requirements of their approach.

"Through this technology, we implemented an NPU with over 60% improved performance compared to the latest GPUs by combining quantisation techniques that reduce memory requirements while maintaining inference accuracy."

The energy-efficient NPU technology developed by KAIST offers a potential path toward more sustainable AI operations. With 44% lower power consumption compared to current GPU solutions, widespread adoption could significantly

reduce the carbon footprint of AI cloud services. However, the technology's real-world impact will depend on several factors, including manufacturing scalability, cost-effectiveness, and industry adoption rates.

The timing of this energy-efficient NPU technology breakthrough is particularly relevant as AI companies face increasing pressure to balance performance with sustainability. The current GPU-dominated market has created supply chain constraints and elevated costs, making alternative solutions increasingly attractive.

Professor Park noted that the technology "has demonstrated the possibility of implementing high-performance, low-power infrastructure specialised for generative AI, and is expected to play a key role not only in AI cloud data centres but also in the AI transformation (AX) environment represented by dynamic, executable AI such as agentic AI."

<https://www.cloudcomputing-news.net/>

Wind turbine blade to capture energy

Researchers have developed a new design platform and a staggering 12-megawatt-class blade to match in an effort to put wind beneath the sails of its domestic production of wind power. The Wind Energy Research Department at the Korea Institute of Energy Research (KIER) set out to increase localization rates for wind turbine components, which stood at around 34%, according to WindTech International. The priority was particularly for larger capacity wind turbines, as Korea still heavily relies on imports to make them.

What resulted was BladeFORGE, a design platform that utilizes both advanced optimization algorithms and artificial intelligence techniques to face the complex technical challenges of building 10-megawatt-plus capacity wind turbines.

The integrated aero structural design system aims to increase design efficiency by addressing aerodynamic and structural factors simultaneously, rather than individually, like in previous manual methods. With this strategy, BladeFORGE cuts optimization time by more than 50%, and the methods have

already received approval in principle from the Korean Register. To match the efficiency of their cutting-edge platform, KIER has also created a research facility at Jeju Global Research Center that will house wind blade design, fabrication, and structural testing all under one roof.

This system and facility are substantial steps in the Republic of Korea's renewable energy efforts, as they both cut manufacturing times domestically and reduce outsourcing of necessary parts. By reducing the need for shipping, in particular, the program reduces the excess use of energy in production that causes the rapid overheating of our planet.

Using BladeFORGE in tandem with this new infrastructure, researchers created a 107-meter (351-foot), 12-megawatt-class wind turbine blade – the first of its size and capacity from Korea to receive Des Norske Veritas design verification. To receive approval from DNV is a major credit to the program, as the provider sets quality and safety standards globally for wind turbine components.

This focus on localizing manufacturing thanks to KIER puts the country's renewable industry at a new level, as it reflects a commitment to the environment at all stages of production, not just in energy production.

<https://www.yahoo.com/>

UNITED ARAB EMIRATES

AI-powered digital twins for clean energy

Researchers at the University of Sharjah have recently conducted an in-depth study looking into how AI-powered digital twins could reshape the future of clean energy. Their findings reveal both the strong potential and ongoing challenges of applying this advanced technology to real-world energy systems.

Digital twins simulate physical systems in real time, allowing engineers to monitor performance, detect issues, and optimise operations without interrupting energy production. In industries like manufacturing and transportation, they've already improved efficiency and reduced costs.

In renewable energy, digital twins can be used to model complicated systems like wind farms, solar arrays, geothermal wells, hydroelectric dams, and biomass supply chains. These simulations can help operators fine-tune performance, predict maintenance needs, and design more intelligent infrastructure.

The researchers conducted a complete review of scientific literature using AI-driven text mining techniques. This allowed them to analyse large amounts of data and identify gaps in existing digital twin applications across renewable energy systems.

Their findings suggest that there is still a need for better data collection methods, more sophisticated modelling techniques, and expanded computational resources. These improvements could improve the accuracy, adaptability, and usefulness of digital twins in energy production and management.

To guide future innovation, the team has proposed a roadmap for overcoming current limitations. Their recommendations hope to strengthen the reliability of digital twins and make them a cornerstone technology in the global effort to reduce carbon emissions.

<https://www.openaccessgovernment.org/>

EUROPE

GERMANY

Method to reduce AI energy consumption

AI applications such as large language models (LLMs) have become an integral part of our everyday lives. The required computing, storage, and transmission capacities are provided by data centers that consume vast amounts of energy. In Germany alone, this amounted to around 16 billion kWh in 2020, or around 1% of the country's total energy consumption. For 2025, this figure is expected to increase to 22 billion kWh.

The arrival of more complex AI applications in the coming years will substantially increase the demands on data center capacity. To counteract this trend, researchers have developed a training

method that is 100 times faster while achieving accuracy comparable to existing procedures.

The functioning of neural networks, which are used in AI for such tasks as image recognition or language processing, is inspired by the way the human brain works. These networks consist of interconnected nodes called artificial neurons. The input signals are weighted with certain parameters and then summed up. If a defined threshold is exceeded, the signal is passed on to the next node.

Felix Dietrich, a professor of Physics-enhanced Machine Learning, and his team have developed a new method. Instead of iteratively determining the parameters between the nodes, their approach uses probabilities. Their probabilistic method is based on the targeted use of values at critical locations in the training data where large and rapid changes in values are taking place. The objective of the current study is to use this approach to acquire energy-conserving dynamic systems from the data. Such systems change over the course of time in accordance with certain rules and are found in climate models and in financial markets, for example.

<https://www.tum.de/>

SWITZERLAND

AI writing climate-friendly cement recipes

When cement is mixed with water, sand, and gravel, it becomes concrete, the most widely used building material in the world. However, the production of cement releases large amounts of carbon dioxide. Researchers at the Paul Scherrer Institute PSI are using artificial intelligence and computational modelling to develop alternative formulations that should be more climate friendly. The Researchers have developed an AI-based model that helps to accelerate the discovery of new cement formulations that could yield the same material quality with a better carbon footprint.

One promising strategy for reducing emissions is to modify the cement recipe itself - replacing some of the clinker with alternative cementitious materials.

That is exactly what an interdisciplinary team in the Laboratory for Waste Management in PSI's Center for Nuclear Engineering and Sciences has been investigating. Instead of relying solely on time-consuming experiments or complex simulations, the researchers developed a modelling approach based on machine learning. "This allows us to simulate and optimise cement formulations so that they emit significantly less CO2 while maintaining the same high level of mechanical performance," explains mathematician Romana Boiger, first author of the study. "Instead of testing thousands of variations in the lab, we can use our model to generate practical recipe suggestions within seconds - it's like having a digital cookbook for climate-friendly cement." With their novel approach, the researchers were able to selectively filter out those cement formulations that could meet the desired criteria.

The researchers at PSI also made use of an artificial neural network. They themselves generated the data required for training: "With the help of the open-source thermodynamic modelling software GEMS, developed at PSI, we calculated - for various cement formulations - which minerals form during hardening and which geochemical processes take place," explains Nikolaos Prasianakis. By combining these results with experimental data and mechanical models, the researchers were able to derive a reliable indicator for mechanical properties - and thus for the material quality of the cement. For each component used, they also applied a corresponding CO2 factor, a specific emission value that made it possible to determine the total CO2 emissions. "That was a very complex and computationally intensive modelling exercise," the scientist says.

"Instead of seconds or minutes, the trained neural network can now calculate mechanical properties for an arbitrary cement recipe in milliseconds - that is, around a thousand times faster than with traditional modelling," Boiger explains.

To find the solution, the team integrated in the workflow an additional AI technology, the so-called genetic algorithms - computer-assisted methods inspired

by natural selection. This enabled them to selectively identify formulations that ideally combine the two target variables. The advantage of this "reverse approach": You no longer have to blindly test countless recipes and then evaluate their resulting properties; instead, you can specifically search for those that meet specific desired criteria - in this case, maximum mechanical properties with minimum CO2 emissions.

<https://www.sciencedaily.com/>

THE NETHERLANDS

AI to face fusion challenges

Experiments produce massive amounts of data. The behavior of plasma, an extremely hot, electrically charged gas, is notoriously difficult to predict. This complexity makes it hard to run experiments efficiently, interpret results, or prevent sudden failures that can damage equipment. To help tackle this problem, PhD researcher Yoeri Poels developed smart data-driven tools that assist scientists in analyzing and controlling fusion experiments. His work supports the safe and efficient development of fusion energy, a clean and potentially limitless power source for the future.

Fusion experiments often use a device called a tokamak, which uses powerful magnetic fields to hold extremely hot plasma in place and allow fusion reactions to happen safely. By analyzing large amounts of data from tokamak fusion experiments, Poels used machine learning to create faster simulation models that save time. He developed more robust monitoring systems capable of detecting subtle changes in the plasma as they happen.

Additionally, he introduced new methods to recognize and study dangerous plasma instabilities, helping scientists prevent equipment damage and improve control within the tokamak. These innovations combined support safer and more efficient fusion energy research.

In his thesis, Poels explored how artificial intelligence, specifically machine learning, can support fusion research.

He developed new methods to address three important challenges:

- **Faster plasma simulations**
Fusion experiments often rely on detailed computer simulations, but these can take a long time to run. Poels created fast data-based simulation tools that learn from past results. These tools are not meant to fully replace traditional simulations, but they can support quicker studies when time or computing power is limited.
- **Monitoring energy performance**
Keeping energy well confined inside the plasma is essential for good fusion results. Poels built a tool that automatically detects how well the plasma is performing, even if some measurements are missing or faulty. It also tells scientists how confident it is in its predictions, which is important for real-time decision making.
- **Understanding disruptions**
Sometimes fusion experiments end in sudden disruptions that can damage equipment. These events are still not well understood. Poels used machine learning to find simplified patterns in large datasets, helping researchers better spot warning signs and analyze past disruptions.

<https://www.tue.nl/>

UK

AI to predict solar energy output

Researchers at the University of Nottingham have created an AI model that allows them to accurately predict the amount of solar energy that can be created in different climates, making grid integration easier in the UK. Solar forecasting, and the ability to predict how much sunlight a certain area might receive, has therefore become more important, prompting researchers in the Faculty of Engineering to find new ways of making this process more reliable.

As a novel approach, researchers have used very-short-term (VST) solar energy forecasting, using ground-based fish-eye images, which has proven effective in predicting rapid and accurate changes in solar irradiance, especially for fast-changing local cloud movements.

To address varied geographical and climatic conditions, the researchers showed that a model initially trained in California's sunny climate can effectively predict solar output in Nottingham, known for its humid and rainy conditions. The approach significantly cut down the amount of local data needed to make accurate forecasts — from four months' worth to just two weeks.

<https://www.russellgroup.ac.uk/>

NORTH AMERICA

USA

AI to improve the efficiency of battery diagnostics

National Renewable Energy Laboratory (NREL) researchers have developed and demonstrated a groundbreaking physics-informed neural network (PINN) model that can predict battery health nearly 1,000 times faster than traditional models.

NREL's PINN replaces the traditional, resource-intensive battery physics model with a powerful artificial intelligence approach that mimics the interconnected neurons of our brains to analyze nonlinear, complex datasets. This deep learning process can enhance battery health diagnostics by quantifying physical degradation mechanisms and pave the way for more efficient, scalable approaches to manage battery aging.

NREL researchers have created a vast array of battery lifespan models to diagnose battery health, predict battery degradation, and optimize battery designs. For years, the team has been on the cutting edge of physics-based machine learning techniques to optimize predictive modelling for advanced battery research.

Two such models, the Single-Particle Model (SPM) and the Pseudo-2D Model (P2D), are widely used and accepted approaches to providing a window into how a battery's internal health parameters—such as electrode inventory and kinetics, Li-ion inventory, and Li transport paths—evolve over time. However, directly using these models is an intensive process that requires massive amounts

of computation and limits their ability to offer rapid diagnostics.

The NREL-developed PINN surrogate combines the predictive power of artificial intelligence with the rigor of physics-based modelling. The resulting two-part study published in the *Journal of Energy Storage* demonstrates how researchers trained and tested the PINN surrogate using conventional SPM and P2D models. This multifaceted approach allowed NREL researchers to train the PINN surrogate on a wide range of internal battery properties. The resulting open-source model offers critical insights into changes that occur during battery aging, helping quickly estimate how long a battery might last in a different setting.

What makes this development especially revolutionary in battery research is the integration of physics-informed principles into neural networks. Traditional neural networks are data-driven models that excel at pattern recognition but often lack the ability to enforce physical laws, which are crucial for accurately simulating battery behaviour.

The success of NREL's PINN surrogate offers wide-ranging implications. For battery diagnostics, the PINN surrogate can provide rapid state-of-health predictions, allowing for faster decision-making across battery applications. By drastically lowering the computational barriers to battery diagnostics, the PINN surrogate model paves the way for widespread, scalable, and efficient energy storage management—helping ensure energy is available when and where it is needed.

Currently, researchers are working to transition the PINN surrogate from controlled simulations to real-world data validation, using batteries cycled within NREL's laboratories. By bridging this gap, researchers hope to deploy PINN-based diagnostics across a wide range of battery systems, enhancing battery performance monitoring and extending lifespans. Future research will focus on refining the PINN model to handle highly dimensional problems, allowing it to predict a broader array of internal battery parameters with increased precision. This means creating models that can both respond to diverse current loads and scale

effectively to future battery designs and usage patterns.

<https://www.nrel.gov/>

AI framework to optimize battery electrolytes

Identifying new, high-performance electrolytes remains a major challenge in developing next-generation batteries for electric vehicles, consumer electronics, and grid-scale energy storage. The most stable electrolytes are not always the most conductive, and the most efficient batteries are not always the most stable, highlighting the complex trade-offs involved.

By analyzing data from 250 research papers covering the history of lithium-ion battery research, the team used AI to calculate an “eScore” for different molecules. This score balances three key properties—ionic conductivity, oxidative stability, and Coulombic efficiency—to highlight top-performing candidates. The researchers validated their AI-driven method by identifying a molecule with performance comparable to today’s leading electrolytes. This achievement represents a significant step forward in a field that has traditionally relied on trial and error.

Artificial intelligence is helping scientists streamline the search for better battery materials by identifying the most promising candidates for lab testing, saving time, energy, and resources. At the University of Chicago’s Pritzker School of Molecular Engineering (PME), researchers are already applying AI to speed up advances in cancer therapies, water purification, quantum materials, and more.

When it comes to battery research, the challenge is immense. The number of possible electrolyte molecules is estimated to be as high as 10^{60} , which is far too many to explore through traditional methods. AI offers a way to narrow this vast field to the most likely candidates.

Amanchukwu compares the use of AI in battery research to how streaming services recommend music. Think of each person’s music taste as their own “eScore.” The current AI can scan a playlist and predict which songs someone might like. The next step is building AI

that can generate an entirely new playlist based on those preferences.

Ultimately, Amanchukwu’s lab is aiming for something even more advanced: AI that cannot just select or predict, but actually design new molecules from scratch that meet all the required performance criteria, like composing new music instead of just recommending it. To support this work, Amanchukwu received a Google Research Scholar Award last year to help move closer to that goal of truly generative electrolyte AI. The team found that the AI performed well when predicting molecules similar to those it had already encountered. However, it was less effective when presented with unfamiliar compounds—a limitation they now aim to overcome in the next phase of their work.

<https://www.azom.com/>

Protecting the grid with AI

The electric grid powers everything from traffic lights to pharmacy fridges. However, it regularly faces threats from severe storms and advanced attackers. Sandia researchers have developed brain-inspired AI algorithms that detect physical problems and cyberattacks both at the same time within the grid. And this neural-network AI can run on inexpensive single-board computers or existing smart grid devices.

As the nation adds more smart controls and devices to the grid, it becomes more flexible and autonomous, but also more vulnerable to cyberattacks and cyber-physical attacks. Cyber-physical attacks use communications networks or other cyber systems to disrupt or control a physical system such as the electric grid. Potentially vulnerable equipment includes smart inverters that turn the direct current produced by solar panels and wind turbines into the alternating current used by the grid, and network switches that provide secure communication for grid operators, said Adrian Chavez, a cybersecurity expert involved in the project.

Because the neural network can run on single-board computers or existing smart grid devices, it can protect older equipment as well as the latest equipment that lacks only cyber-physical

coordination, Shamina said. “To make the technology more accessible and feasible to deploy, we wanted to make sure our solution was scalable, portable, and cost-efficient,” Adrian said.

The Sandia team collaborated with experts at Texas A&M University to create secure communication methods, particularly between grids owned by different companies, Shamina said. The biggest challenge in detecting cyber-physical attacks is combining the constant stream of physical data with intermittent packets of cyber data, said Logan Blakely, a computer science expert who led development of the AI components.

Physical data, such as the frequency, voltage, and current of the grid, is reported 60 times a second, while cyber data, such as other traffic on the network, is more sporadic, Logan said. The team used data fusion to extract the important signals in the two different kinds of data. The collaborators from Texas A&M University were key to this effort, he added.

Then the team used an autoencoder neural network, which classifies the combined data to determine whether it fits with the pattern of normal behaviour or if there are abnormalities with the cyber data, physical data, or both, Shamina said. For example, an increase in network traffic could indicate a denial-of-service attack, while a false-data-injection attack could include atypical physical and cyber data, Adrian said.

Unlike many other kinds of AI, autoencoder neural networks do not need to be trained on data labelled with every type of issue that may show up, Logan said. Instead, the network only needs copious amounts of data from normal operations for training. The use of an autoencoder neural network makes the package pretty much plug and play, Shamina added.

Once the team constructed the autoencoder neural network, they put it to the test in three different ways. First, they tested the autoencoder in an emulation environment, which includes computer models of the communication-and-control system used to monitor the grid and a physics-based model of the grid itself, Shamina said. The team used this

environment to model a variety of cyberattacks or physical disruptions, and to provide normal operational data for the AI to train on. The collaborators from Texas A&M University assisted with the emulation testing.

Then the team incorporated the autoencoder onto single-board computer prototypes that were tested in a hardware-in-the-loop environment, Shamina said. In hardware-in-the-loop testing, researchers connect a real piece of hardware to software that simulates various attack scenarios or disruptions. When the autoencoder is on a single-board computer, it can read the data and implement the algorithms faster than a virtual implementation of the autoencoder can in an emulation environment, Adrian said. Generally, hardware implementations are a hundred or thousand times faster than software implementations, he added.

The team is working with Sierra Nevada Corp. to test how Sandia's autoencoder AI works on the company's existing cybersecurity device called Binary Armor, Shamina said. The team is testing both formats — single-board prototypes interfaced with the grid and the AI package on existing devices — in the real world at the Public Service Co. of New Mexico's Prosperity solar farm as part of a Cooperative Research and Development Agreement, Shamina said. These tests began last summer, Adrian said.

The team also worked with PNM early in the project to learn what AI design might be most useful for grid operators. It was during conversations with PNM staff that the Sandia team identified the need to connect cyber-defenders with system operators rapidly and automatically.

This project built off and expanded upon a previous R&D 100 Award-winning

project called the Proactive Intrusion Detection and Mitigation System, which focused on detecting and responding to cyber intrusions in smart inverters on solar panels, Shamina said. The team is also expanding upon the autoencoder AI in similar projects, she added.

The team filed a patent on the autoencoder AI and is looking for corporate partners to deploy and hone the technology in the real world, Shamina said. With a bit more work, the autoencoder could be used to protect other critical infrastructure systems such as water and natural gas distribution systems, factories, and even data centers, Adrian said.

The project is funded by Sandia's Laboratory Directed Research and Development program.

<https://www.sandia.gov/>