The AI Scientist

Physics Informed Neural Networks

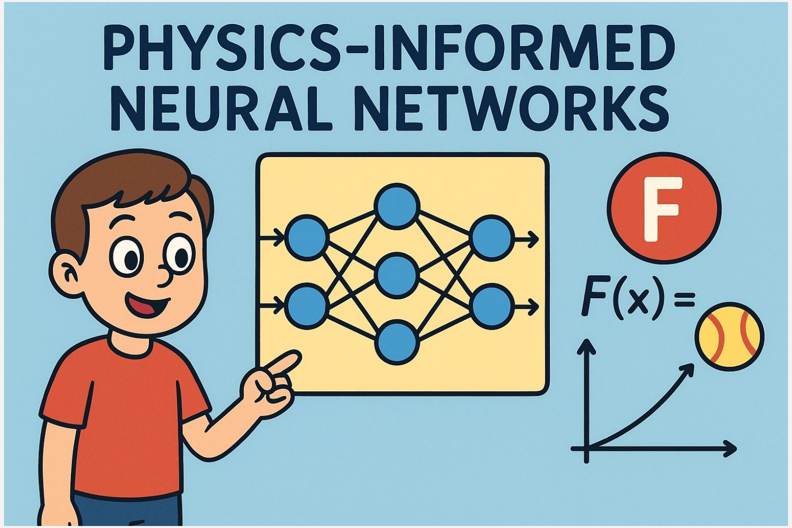

Outline

Core Idea

Motivating Examples

Theory (Recipe)

Applications

Learning Outcome

By the end of this session, you will understand:

- The Core Idea: What a Physics-Informed Neural Network (PINN) is?
- The Key Difference: How it's better than a data-only AI:
  - Data AI: Memorizes patterns.
  - PINN: Discovers rules and understands the *"why."*
- How It Works?
- The Power of PINNs where data is scarce or expensive.
- Applications.

We all know that AI is really good at learning from information, which we call 'data'.

You've probably also heard of AIs like ChatGPT. It reads millions of books and websites to learn how to use words and answer questions. Its 'data' is all that text.

But what if we wanted an AI to understand the *physical world*?

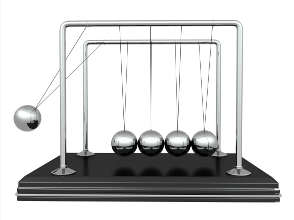

Image Credit: adventtr / Getty Images

To predict how things move, heat up, or flow?

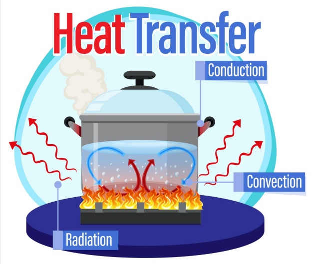

Image Credit: blueringmedia / Getty Images

We could just show it a lot of data, but that's expensive and sometimes impossible (like predicting the weather on Mars!).

Image Credit: Mark Stevenson / Getty Images

Instead, we can give it a secret weapon: the cheat sheet of the universe—the laws of physics!

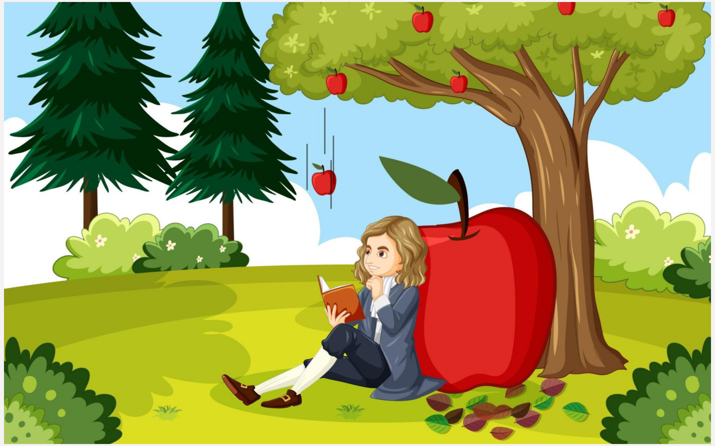

Image Credit: blueringmedia / Getty Images

This super-smart AI that uses both data and the rules of physics is called a Physics-Informed Neural Network (PINN). But let's just call it our "AI Scientist."

How it Works: The "Guess, Check, and Learn" Game

The Guess: The AI Scientist looks at a situation (e.g., a swinging pendulum) and makes a wild guess about what will happen next. At first, its guesses will be terrible!

Image Credit: Rubberball/Mike Kemp / Getty Images

The "Physics Check": This is the special part. We don't just tell the AI it's wrong. We ask it: *"Does your guess follow Newton's Laws of motion? If you calculate the forces, does it make sense?"* We give it a penalty every time its guess breaks a physics law.

Image Credit: rudall30 / Getty Images

The Learn: The AI hates getting penalties. So, it slowly adjusts its guesses to not only fit the little data it has but also to obey the laws of physics. It becomes a more disciplined scientist.

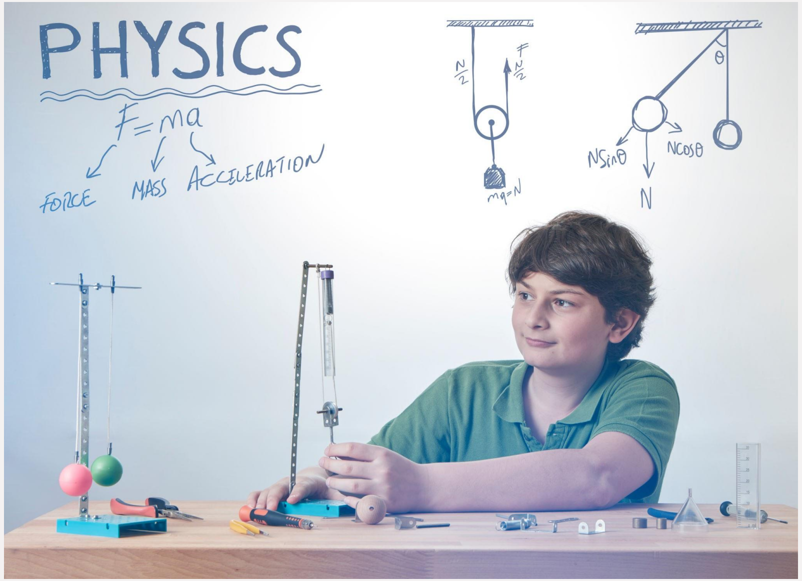

Image Credit: xefstock / Getty Images

Example 1: The Swinging Pendulum (Motion)

The Problem: We want to predict exactly where this pendulum will be at any moment in the future.

Image Credit: Rubberball/Mike Kemp / Getty Images

The Old Way (Just Data): We film the pendulum for an hour, noting its position every second. We feed this thousands of data points to a normal AI. It might learn the pattern, but if we change the weight of the bob, its predictions will be useless. It didn't learn why it moves, just how it moved in that one experiment.

The "AI Scientist" Way (PINN):

1. We give it just a few data points (where the pendulum was at the start and a couple of swings).

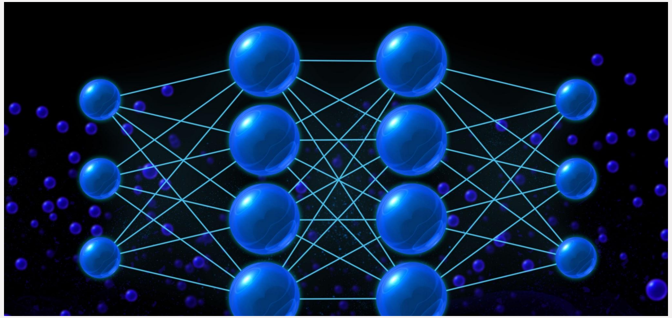

Image Credit: Rubberball/Mike Kemp / Getty Images

2. We give it the "**Physics Cheat Sheet**": Newton's Second Law (`F = m*a`). We tell the AI: "Your predictions for position must satisfy this equation for gravity and tension.

Image Credit: Rubberball/Mike Kemp / Getty Images

3. The AI makes guesses, checks them against `F = m*a`, gets penalized for bad guesses, and learns to make perfect predictions that work for *any* weight or string length! It has discovered the true law of the motion.

Example 2: The Cooling Mug of Cocoa (Heat)

The Problem: How long do I have to wait before my hot cocoa is the perfect drinking temperature?

The Old Way (Just Data): I measure the temperature every minute for 30 minutes. I plot the points and draw a curve. This works for my mug, in my room. But if I go outside on a cold day, my curve is wrong.

Image Credit: valentinarr / Getty Images

The "AI Scientist" Way (PINN):

1. We give it a few temperature measurements (e.g., at 0, 5, and 10 minutes).

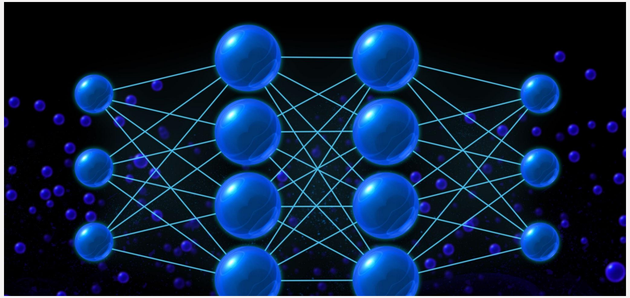

Image Credit: Rubberball/Mike Kemp / Getty Images

2. We give it the **"Physics Cheat Sheet"**: Newton's Law of Cooling (The rate of cooling is proportional to the difference between the object's temperature and the room's temperature). Tell the AI: "Your predicted cooling curve must obey this law."

Image Credit: Rubberball/Mike Kemp / Getty Images

3. The AI learns a model that understands the *principle* of cooling. Now, I can tell the AI "the room is colder today," and it will accurately predict how that changes the cooling time! It's smarter because it knows the *why*.

Example 3: Water Ripple in a Pond (Waves)

The Problem: If I drop two stones in a pond, can I predict how their ripples will crash into each other?

The Old Way (Just Data): This is nearly impossible! I would have to place thousands of sensors all over the pond to measure the water height. It's incredibly expensive and messy.

Image Credit: IRA_EVVA / Getty Images

The "AI Scientist" Way (PINN):

1. We show it a short video of the first few ripples starting to spread. That's our only data.

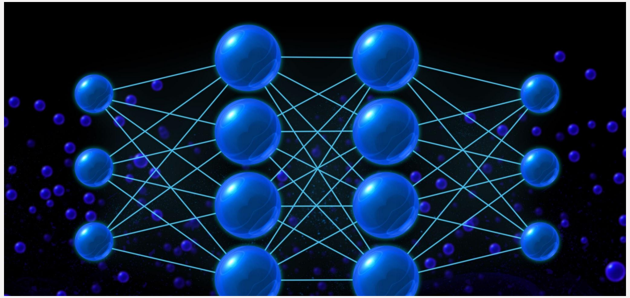

Image Credit: Rubberball/Mike Kemp / Getty Images

2. We give it the **"Physics Cheat Sheet"**: The Wave Equation (a famous rule that describes how all waves behave). We tell the AI: "Your simulation of the future must obey this wave equation."

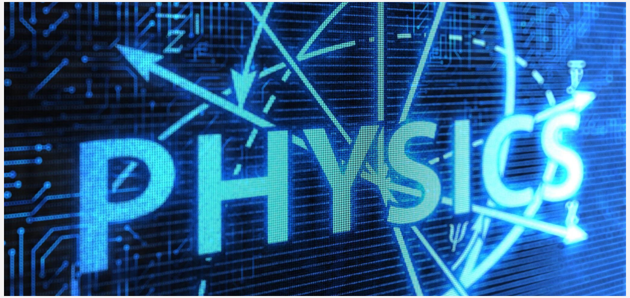

Image Credit: Rubberball/Mike Kemp / Getty Images

3. The AI takes the initial video, uses the wave equation as a guide, and can now simulate on its own how the ripples will propagate and interfere far into the future, far beyond what the video showed! It's using physics to fill in the massive gaps in the data.

The AI Scientist's Recipe

Ingredient 1: The Data (The "Examples")

- **What it is:** This is just a few real-world measurements. It's the "what we know for sure."
- **Analogy:** It's like the few key points you're given to connect in a "connect-the-dots" picture.

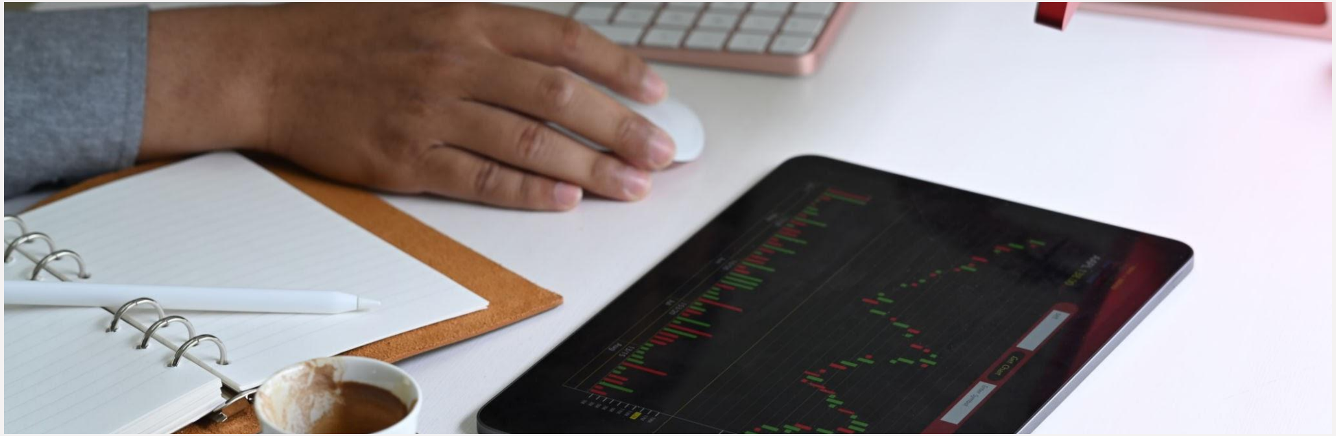

Image Credit: PrathanChorruangsak / Getty Images

Data

Examples:

- For the pendulum: A few (time, position) pairs.
- For the cocoa: A few (time, temperature) pairs.
- For the ripple: A short video of the beginning.

Image Credit: Choys-Design / Getty Images

This gives the AI a starting point, something to anchor its guesses to reality.

Ingredient 2: The Physics Rule (The "Cheat Sheet")

- **What it is:** This is the scientific law, written in the language of math.
- **Analogy:** It's the **rulebook** for the connect-the-dots game. It says, "The line you draw between the dots must be smooth and curvy," not jagged and crazy.

Image Credit: Junior Jap / Getty Images

The Physics Rule

Examples

- For the pendulum: "A rule that calculates how the **push** (force) depends on the **angle**."
- For the cocoa: "A rule that says the **speed of cooling** depends on how much **hotter the drink is than the room**."
- For the ripple: "A rule that describes how a **wave's shape spreads out** over time."

This is what makes the AI a scientist. It's the logic it uses to critique its own work.

Ingredient 3: The Combined "Wrongness" Meter (The "Learning Engine")

What it is: This is the coolest part. The AI has a single goal: be less wrong.

- It has a "Data Wrongness" Score: How far its guess is from the actual data points (the dots).
- It has a "Physics Wrongness" Score: How much its guess breaks the physics rule (the rulebook).

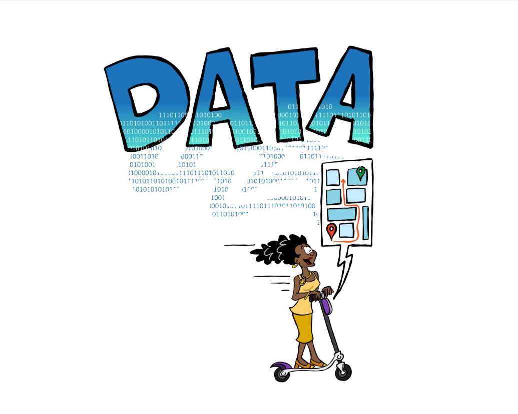

Image Credit: Antonio MEZA / Getty Images

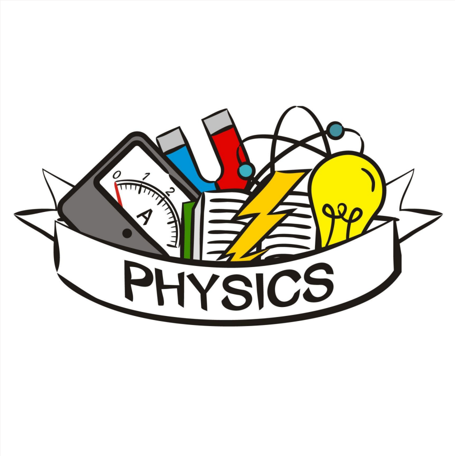

Image Credit: Yulia Andreeva / Getty Images

The Combined "Wrongness" Meter

Analogy: Imagine you're drawing that connect-the-dots picture.

- **Data Wrongness:** How far your pen is from the dots.
- **Physics Wrongness:** How jagged and messy your line is between the dots.
- **Total Goal:** Draw a line that hits all the dots *and* is smooth. You have to balance both!

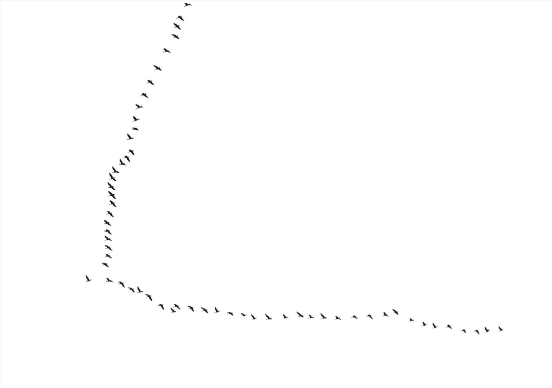

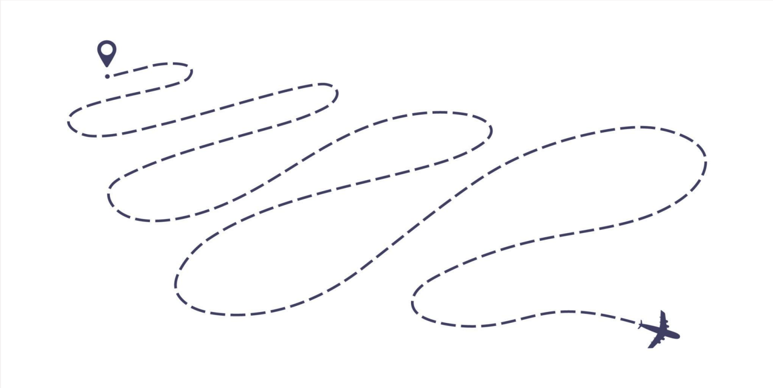

The AI's brain is designed to constantly try new guesses and measure the total wrongness. It keeps adjusting until it finds the answer that makes both the data wrongness and the physics wrongness as small as possible.

Applications

Geophysics & Climate Science

Mapping the Earth's Secrets [1]

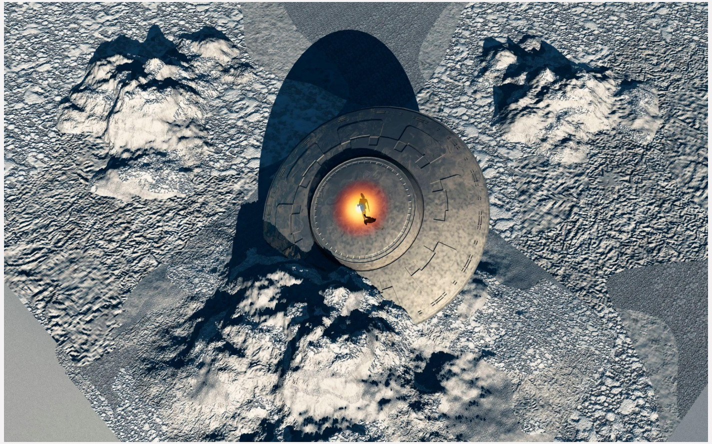

Image Credit: Mark Stevenson/Universal Images Group / Getty Images

Predicting Extreme Weather [2]

Image Credit: Rafael-Chacon-Photography / Getty Images

[1] Wu, X., Ma, J., Si, X., Bi, Z., Yang, J., Gao, H., ... & Zhang, J. (2023). Sensing prior constraints in deep neural networks for solving exploration geophysical problems. Proceedings of the National Academy of Sciences, 120(23), e2219573120.
[2] Kashinath, K., Mustafa, M., Albert, A., Wu, J. L., Jiang, C., Esmaeilzadeh, S., ... & Prabhat, N. (2021). Physics-informed machine learning: case studies for weather and climate modelling. Philosophical Transactions of the Royal Society A, 379(2194), 20200093.

Medicine & Biology

Personalized Blood Flow [1]

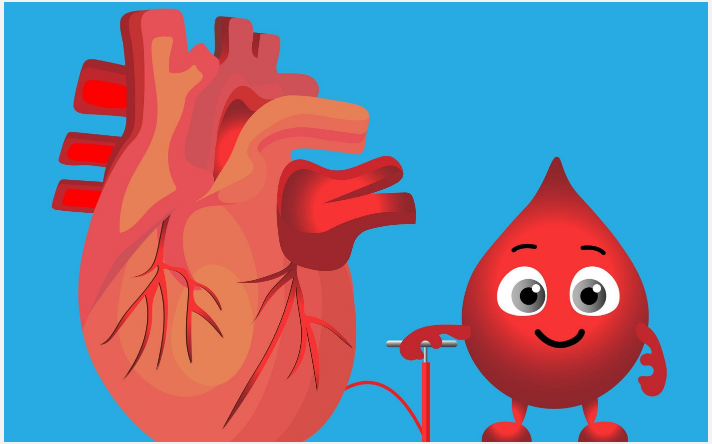

Image Credit: Alfieri / Getty Images

Unveiling the Invisible (Digital Twin of a Heart) [2]

Image Credit: vege / Getty Images

[1] Bhargava, S., & Chamakuri, N. (2024). Enhancing arterial blood flow simulations through physics-informed neural networks. arXiv preprint arXiv:2404.16347.
[2] Kuang, K., Dean, F., Jedlicki, J. B., Ouyang, D., Philippakis, A., Sontag, D. A., & Alaa, A. M. (2024). Non-invasive medical digital twins using physics-informed self-supervised learning. CoRR.

Engineering & Design

The Invisible Aerodynamics of Race Cars [1]

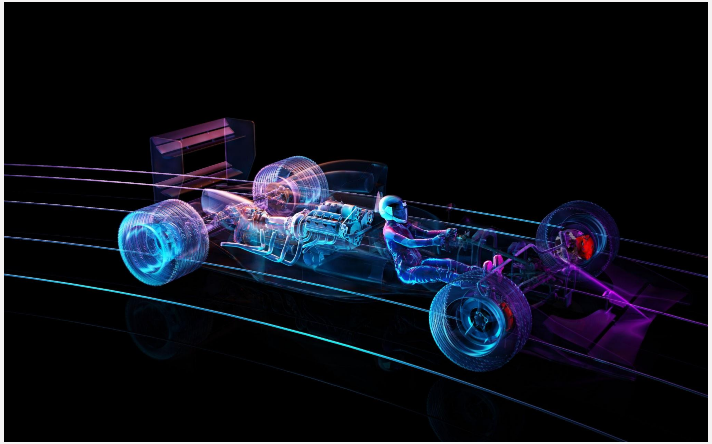

Image Credit: 1971yes / Getty Images

Cracking the Code of Fusion Energy [2]

Image Credit: koto_feja / Getty Images

[1] Shah, N. (2025). Computational Fluid Dynamics Optimization of F1 Front Wing using Physics Informed Neural Networks. arXiv preprint arXiv:2509.01963.
[2] Prantikos, K., Chatzidakis, S., Tsoukalas, L. H., & Heifetz, A. (2023). Physics-informed neural network with transfer learning (TL-PINN) based on domain similarity measure for prediction of nuclear reactor transients. Scientific Reports, 13(1), 16840.

Fundamental Science

Discovering New Laws of Nature

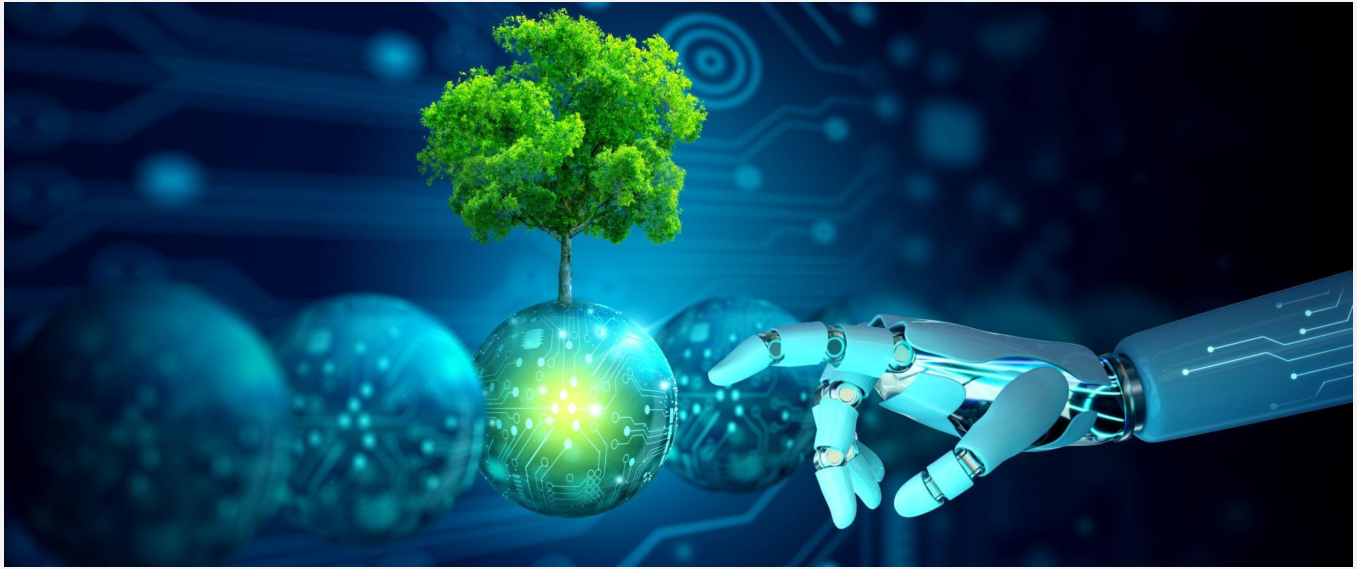

Image Credit: Peach_iStock / Getty Images

[1] Ganga, S., & Uddin, Z. Exploring physics-informed neural networks: From fundamentals to applications in complex systems. arXiv 2024. arXiv preprint arXiv:2410.00422.

Limitations of PINNs [1]

1. Slow Learners:

PINNs take longer to train because they must learn from both **data** and **physics equations** at the same time.

2. Need the Rules First:

They work only if we already know the **physics laws** (like gravity or heat flow). If the rule is unknown, PINNs can't learn it easily.

3. Struggle with Messy Real-World Systems:

When physics is very complex (like turbulent air or weather), PINNs may give rough or unstable results.

4. Hard to Balance:

It's tricky to decide **how much to trust data vs physics**, so sometimes predictions can lean too much toward one side.

[1] Krishnapriyan, A., Gholami, A., Zhe, S., Kirby, R., & Mahoney, M. W. (2021). Characterizing possible failure modes in physics-informed neural networks. Advances in neural information processing systems, 34, 26548-26560.

Image Credit: Suchat longthara / Getty Images