Geminus

Physics-Informed AI

Beyond the Cult of the Data-Driven

Traditional AI comes with a massive data burden, making it challenging to move into industrial production environments. This whitepaper outlines the three big data challenges with AI, and explores a hybrid, physics-informed solution that results in accurate, fast, and explainable models.

AI's Data Problem

Artificial intelligence promises to remake the world. The fuse has been lit, and headlines abound about multi-trillion dollar economic impacts, claims of sentience, and even art creation. In the industrial world, we worry less about robots taking over the world, and more about how AI can enable a leap in the industrial realm. How can AI change the way we design and make products? The way we optimize our processes? And how do we accomplish this while battling the ongoing effects of climate change?

Inherent to the AI craze is a data burden. To create accurate models that support decision-making, there is a constant hunger surrounding data – always more, cleaned, and labeled to ensure that models won't fail when put into action. It also requires an incredible amount of computing power needed to train ever-hungrier AI models. This of course supposes that you have this data, which, in industrial applications, you probably don't. In practice, you're more likely to run into the following scenarios:

01 The Data is Dirty

Any data scientist can tell you how much pain there is in cleaning data before it ever even gets to the model. According to Anaconda's 2021 survey, 39% of data scientist time is spent on data prep and cleaning. Given the importance of data quality on AI results, it's the right place to spend time.

But it does beg the question, what if that time could be reduced or even eliminated? What would your team be able to accomplish with nearly 40% of their time back?

02 The Data is Sparse, Indirect, or Hard to Get

Having insight into a process (like a chemical reaction) or system (like a pressure or a temperature at an arbitrary location) is critical, but direct sensorization is often impossible or costly. Instead, engineers infer behavior, often inaccurately, from the few sensors or data points they do have.

Product designers, on the other hand, are faced with the challenge that their product or process doesn't exist yet, so there aren't any sensors to deploy. Prototyping to get that data must be done judiciously and is often a non-starter due to time and cost constraints.

Both examples are the reason simulation exists. More on that shortly.

03 The Data is Narrow

In industrial applications, data on all potential scenarios is simply not available. For example, if you want to train a model to predict when a motor will fail but the motors don't fail very often, then it will be very difficult to capture the signs of failure with the data you have. Machine learning models are only representative of the data that they've been trained on.

Simulation: Solid, But Not Sexy

Let's step away from AI for a moment and look at a tried-and-true approach used to predict the behavior of complex systems, long before the recent AI craze. Simulation relies on the fact that there are equations (often describing first principles) that reliably predict the impact of changing inputs on a particular output. These tools are ubiquitous, primarily in the fields of product and process design. Once again, we have a powerful tool, though one that still has limitations and challenges:

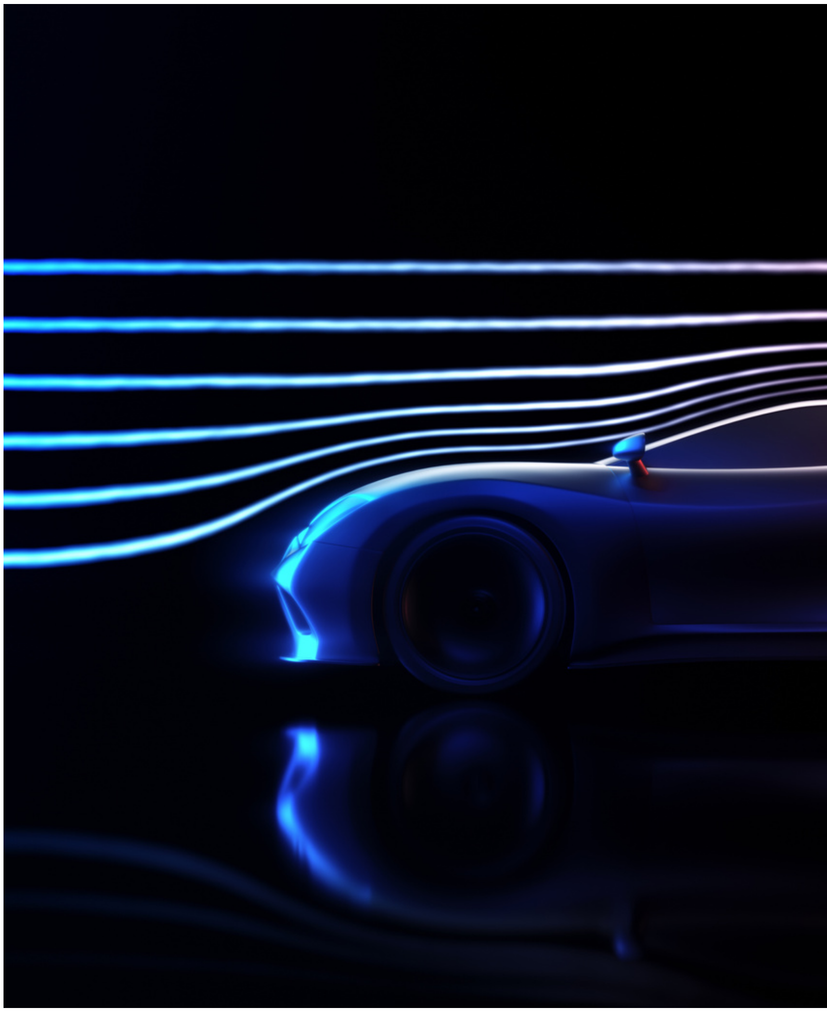

01 Simulation is Slow

The word "slow" here represents a few facets. At the core, it refers to the computational requirements necessary to solve those complex equations. For difficult problems like the complete aerodynamic performance of a car, or the thermodynamics of a natural gas refinery, it can take hours, days, weeks, and even months to compute a single solution. This slows the iterative process, which lengthens design cycles, and makes real-time process optimization impossible.

02 Simulation Doesn't Always Converge

In some cases, you can wait all that time for your answer and find that the answer makes no sense and is unreliable. This is the concept of non-convergence. The net effect of non-convergence is that you cannot always guarantee useful results, which either further lengthens the design cycle or compromises your ability to optimize processes in real-time.

03 Refining Simulation is Hard

For systems with many parameters and independent variables, accurately tuning your simulation to get the most realistic result is hard. In some cases, this becomes more art than science, with experts tweaking knobs based on experience and intuition to arrive at acceptable answers. Once again, this impacts both the result and the time to get that result.

That said, simulation is an immensely valuable tool. In many problems, it is the only way to get insight into how a complex system will behave. This class of data-constrained problems is not a great fit for pure AI, and represents an immense opportunity if some of these simulation issues can be solved.

Joining Forces To Do Better

If we abstract the objective of data-driven AI, it can be summed up as **ensuring reliable data which, in turn, enables more reliable AI models**. To bring it full circle, what if simulation could be the source of that reliable data? This is the field of hybrid modeling, or as we call it, physics-informed AI. At Geminus, we use differentiated methodologies to inject real-world data into simulation to ensure maximum accuracy, then create machine learning models that retain the richness of the simulation but with the robustness and speed of a machine learning model. The upshot of this is we now have models that are 1) accurate 2) fast and 3) explainable. **What can we do with such a model?**

Change the Design Paradigm

We described earlier the inefficient back and forth inherent to the design process. A product designer wants to know the impact of changing process parameters on performance and so submits a request to the modeling team. The modeling team runs the simulation, which may or may not converge, and which can take days or weeks. The product designer decides that the design isn't feasible and wants to explore 5 other iterations, which drags out the design cycle. Imagine a model-driven application that lets you quickly change inputs such as geometry and get the results in a second or less.

Create a Thread Between Design and Operations

Simulation has largely been confined to the design world due to speed issues. A fast and accurate model can now be used to design a complex process and then give real-time insights into future performance for process optimization. Now, the same model can give recommendations on optimal settings to improve yield and maximize energy efficiency due to its predictive power.

Create *Real* Digital Twins

The definition of a digital twin is widely debated. Borrowing from McKinsey, we believe the highest evolution of a digital twin will "use predictions of component failure rates or performance variations to react to changing environments and manipulate the real-world counterpart in a closed-loop setup." Today, most digital twins are glorified visualizations that capture the present, rarely predict the future, and most certainly cannot directly control real live equipment and systems. High-performing models that combine real-world accuracy and real-time speed can enable this vision, which is the promise of physics-informed AI.

Looking through either the AI or simulation lens we find bottlenecks and challenges. Physics-informed AI is a path towards combining the benefits of both, while overcoming their shortcomings. The world of problems governed by physics yet without a feasible path to high-performing models is huge, untapped, and ripe for disruption.

Get in Touch

geminus.ai

info@geminus.ai

geminus.ai