CAE COMPUTATIONAL SCALING CHALLENGES SOLVED BY DATA AND AI-BASED PHYSICS

Sridhar Ravikoti – Director of Product Management, Altair / Joseph Pajot – Director of EDS Product Management, Altair /
September 15, 2025

Introduction

Once regarded as specialized tools, physics simulations have become a standard part of the product design cycle. Structural analysis and fluids simulations are now ubiquitous in industries ranging from automotive and aerospace to consumer products and more. Simulation-driven design has replaced most physical prototyping as it enables early and frequent assessments of a product's design space. In fact, many workflows relegate prototypes to the validation of simulations and the verification of near-final designs due to their associated costs, time, and resources. Thus, simulations improve product KPI, time-to-market, and costs.

As products become more complex, however, simulations have become a new bottleneck for product design cycles. Valuable insights upfront mean teams can collaborate rapidly to evaluate and create designs. Unfortunately, the computational power needed to produce these complex simulations with reasonable speed, accuracy, and cost is becoming a challenge for large organizations — let alone small-to-medium businesses. This reduces collaboration and effective design space exploration.

Fortunately, simulation-heavy, design processes produce a lot of data that can be used as a catalyst to nullify the bottleneck challenges of complex simulations. Wise organizations that recognize the value of simulation data are using it to inform future designs via artificial intelligence (AI). For example, tools like Altair® PhysicsAI™ use this data to train models that return results orders-of-magnitude faster with all the accuracy of full simulations. When PhysicsAI is paired with tools like Altair One®, AI models can be trained quickly mitigating the major bottleneck associated with AI tools.

Challenges

Organizations in the initial stages of adopting simulation-driven design often accomplish their goals on workstations. To remain relevant in the market, however, engineers must continually produce the next, best iteration of their products. Thus, success broadens a product's scope, complexity, and design space. Traditionally, more powerful computing resources were purchased to bear the computational loads needed to design these complex products. For example, enterprise level teams are known to purchase on-site, high-performance computing (HPC) clusters once simulations become too complex for workstations.

Increasing complexity is not the only challenge putting product development teams into crunch mode. Customers now also expect frequent, often yearly, updates to products. This forces engineers to hasten design cycles as their work grows in complexity. That has incentivized organizations to adopt cloud solutions with elastic HPC resources. The idea is to rent the resources needed to compute simulations as they are needed. This avoids the requirement to frequently purchase or update HPC resources.

There are limitations to how simulations scale with computational resources; some organizations are seeing cloud-based simulations become a development bottleneck. This is when AI-based simulations become essential. For example, PhysicsAI models can produce accurate results in seconds — often representing a 1000x speed-up compared to numerical simulations. This substantial time savings provides the answer to today's furious pace of product development. The biggest challenge of AI is that a significant amount of data and computational resources are required to train models. By using cloud computing tools, like Altair One, and an organization's historical simulation data, this challenge can be mitigated.

PhysicsAI: Design Insights Up to 1000x Faster than Traditional Solver Simulations

PhysicsAI models deliver fast physics predictions by learning from historical simulation data. Unlike traditional surrogate models — like reduced order models (ROMs) — PhysicsAI models are not limited to parametric studies. The models operate directly on meshes and CAD geometry. In essence, PhysicsAI models identify the relationship between shapes, performance, and physics by transforming historical data into insights. They then streamline the early design process by taking a fraction of the time to produce results when compared to traditional solvers. PhysicsAI puts innovative and easy-to-use geometric deep learning in the hands of engineers, empowering them to explore more concepts and make better design decisions.

Embedded in Altair® HyperWorks®, PhysicsAI is compatible with various CAE tools on the platform. It is also available as a standalone product that can interface with various third-party software via API tools. PhysicsAI models can also improve product development budgets as their speed — on even basic workstations — reduce the need for HPC environments. Finally, PhysicsAI fits into any organization as it can be deployed on workstations, HPC, and cloud environments.

The general steps to train models in PhysicsAI include:

1. **Collect Data:** Engineers identify data sources (like historical simulation data), curate them for quality, and continue to maintain the data over time. Overlooking this fundamental step is frequently the biggest impediment to success.
2. **Train a Model:** PhysicsAI's guided workflows help engineers set up model training processes and interpret the results. With it, engineers with minimal AI training can build advanced geometric deep learning models.
3. **Assess Quality:** Engineers verify and validate the model on a dataset that is separate from the training data. This ensures predictions are generalized to new, previously unseen, data and are not simply memorizing training data.
4. **Iterate Steps 1 to 3:** Engineers retrain the model based on new data and the quality of the current model. They then iterate continuously until they have a model that accurately predicts simulation results.
5. **Predict Results:** Engineers now use the model to predict physics behavior in a fraction of the time of conventional simulations. The savings will improve product results for engineering design teams.

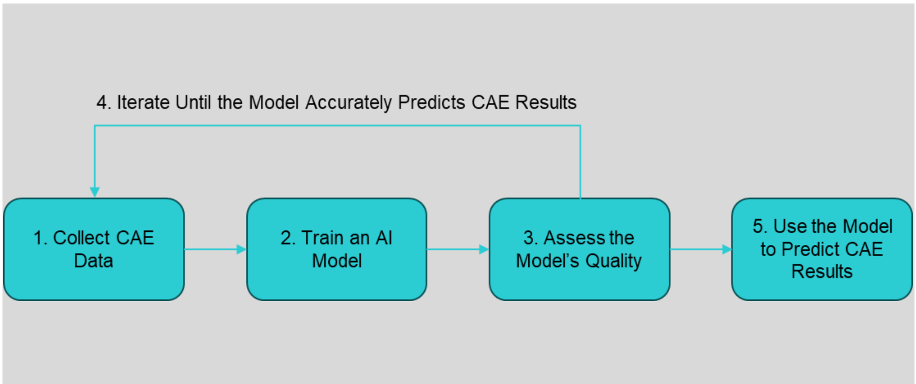

Figure 1 – The general workflow of a PhysicsAI project.

The most computationally intensive step of this process is the training phase. Training times vary from hours to days, depending on the complexity of the problem. The time is well spent as the models produced generate results orders-of-magnitudes faster than traditional simulations — providing an answer to today's rapid pace of product development. Additionally, this time investment can be significantly reduced by using cloud computing tools, like those available on Altair One.

Altair One: The Gateway to Secure, Centralized, AI-based Simulation Process and Data Management

To unlock the full potential of AI, companies must address gaps in skills, data management, compute infrastructure, and application integration. Altair One is built to overcome these obstacles and accelerate the journey toward enterprise-scale AI adoption; its user experience and user interface (UI/UX) make AI workflows accessible to various skill levels. It also combines the foundational pillars of physics-based simulation, data science, and HPC.

Bridging the Skills Gap

A major challenge in adopting AI is the shortage of expertise among engineering teams. Altair One addresses this through seamless integrations into PhysicsAI's guided, engineering workflows. These solutions combine to help engineers apply AI techniques, such as machine learning frameworks, with minimal training.

Breaking Down Data Silos

Data management is the foundation of any AI initiative. Altair One's data-centric architecture brings that management and coherence — connecting simulations, models, and results across departments. Through built-in workflows and metadata extraction, Altair One makes it possible to version, search, and visualize information across the product lifecycle. It establishes a traceable digital thread, ensuring the right figures are accessible at every stage (from model training to validation).

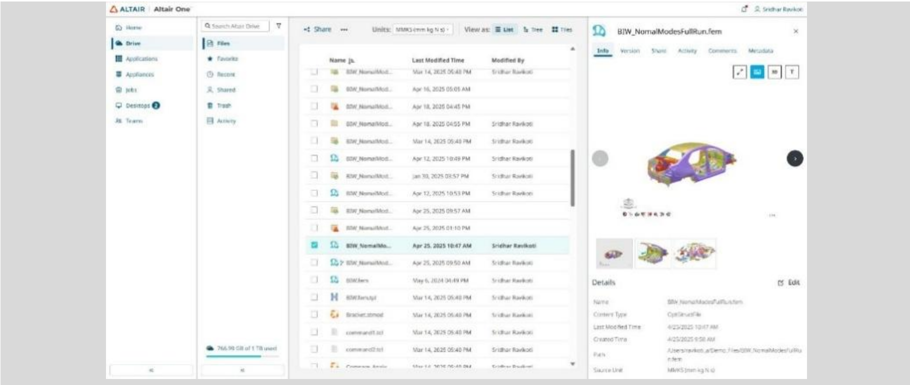

Figure 2 – A screen capture of the Altair One drive. Engineers can use it to search metadata and analyze results in a web-native, 3D viewer. This ensures AI models are trained on quality data.

Compute at Enterprise Scale

AI and simulation workloads demand robust compute infrastructure; Altair One offers on-demand compute capabilities. With a few clicks, users can spin up HPC clusters. This model hides the complexities of software licensing, job scheduling, security, networking, and configuration. It also enables engineers to focus on solving problems rather than managing infrastructure.

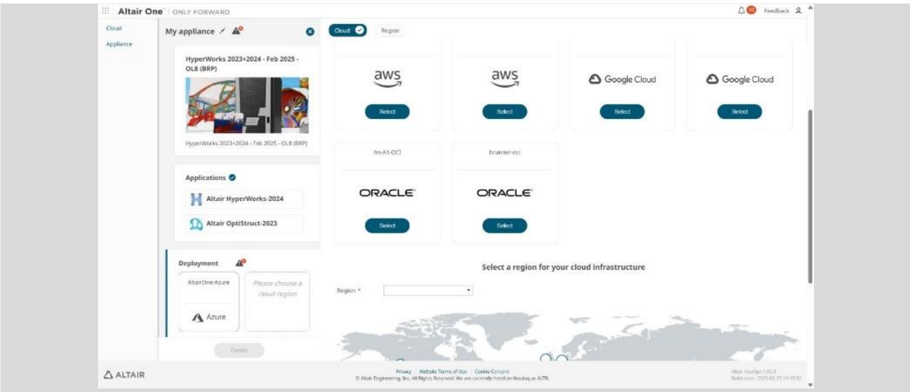

Figure 3 – A screen capture of Altair One's self-service compute infrastructure. Administrators use it to manage licensing, networking, security, and the configurations of provisioned HPC systems.

Application Integration Without Complexity

Modern AI workflows require access to a wide swath of applications. Altair One streamlines this with its fully integrated application ecosystem. These tools are provisioned alongside compute clusters to ensure a consistent user experience. Granular access controls also ensure the right users, and teams, access the tools they need when they need them.

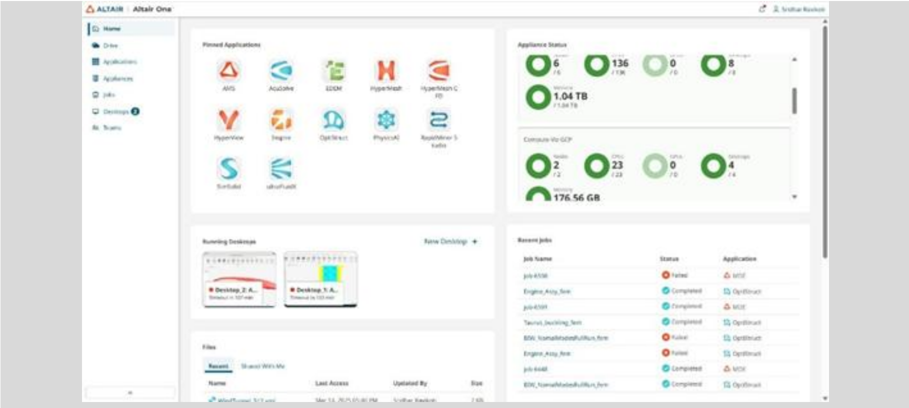

Figure 4 – A screen capture of Altair One's user dashboard. Engineers have direct access to integrate applications, storage, and compute clusters.

All Offerings of Altair One Include an AI-Powered Engineering Ecosystem

Altair One is more than a platform — it is an ecosystem built to support the next era of intelligent engineering. It breaks down silos, consolidates tools, and democratizes access to AI-powered workflows. It is available in three delivery modes:

- Altair One Cloud: a software as a service (SaaS) offering hosted by Altair.
- Altair One Enterprise: a deployment within a customer's infrastructure (on-premises or their own cloud).
- Altair® Unlimited™: a fully managed hardware and software bundle, offered as a physical or virtual appliance.

Experiments and Results

Due to its ubiquity, a standardized automotive crash simulation was selected to demonstrate the synergies of PhysicsAI's fast simulations and Altair One's scalable HPC. These simulations are mandated by regulations, so they are performed frequently — producing a lot of data with varying geometries and initial conditions. They are also familiar, even to those outside the automotive industry, making this assessment accessible to a wide engineering audience. The specific simulation selected assesses the impact between a vehicle and a pedestrian's head (Figure 5).

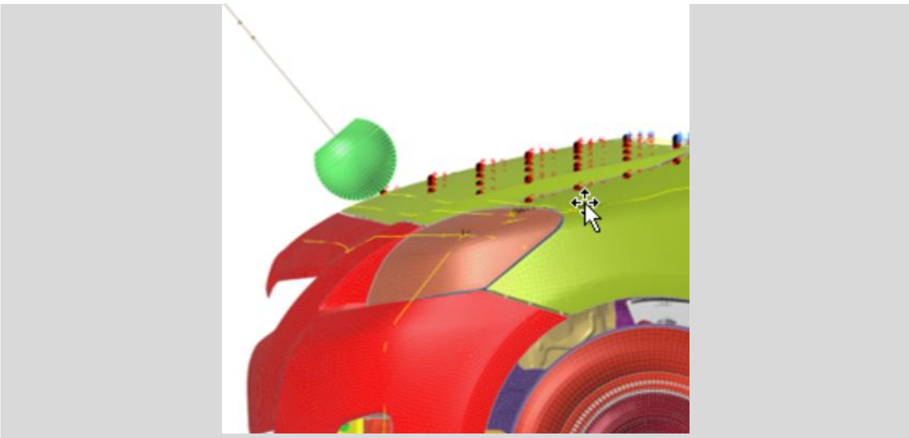

Figure 5 – A finite-element, transient simulation of a vehicle impacting a head.

Comparing the AI and Simulation Results

PhysicsAI trained the models on the datasets of 160 existing head impact simulations and 3000 epochs. The data contained simulations with varying head-form impact locations, hood thicknesses, and initial conditions. After training with a conventional 80/20 train/test split, the PhysicsAI model predicted head-form impact performance solely from a vehicle's geometry. Furthermore, the prediction included both the complete displacement history (as a field quantity) and the peak acceleration (as a scalar quantity).

Figure 6 illustrates a comparison of the displacement contour between the AI prediction and the simulation. Its plots show excellent correlation between the models. A standard workstation laptop returned the AI-based results in about six seconds, and the traditional simulation results in about 30 minutes. Coupled with the high accuracy between the models, this represents a 300x speed boost.

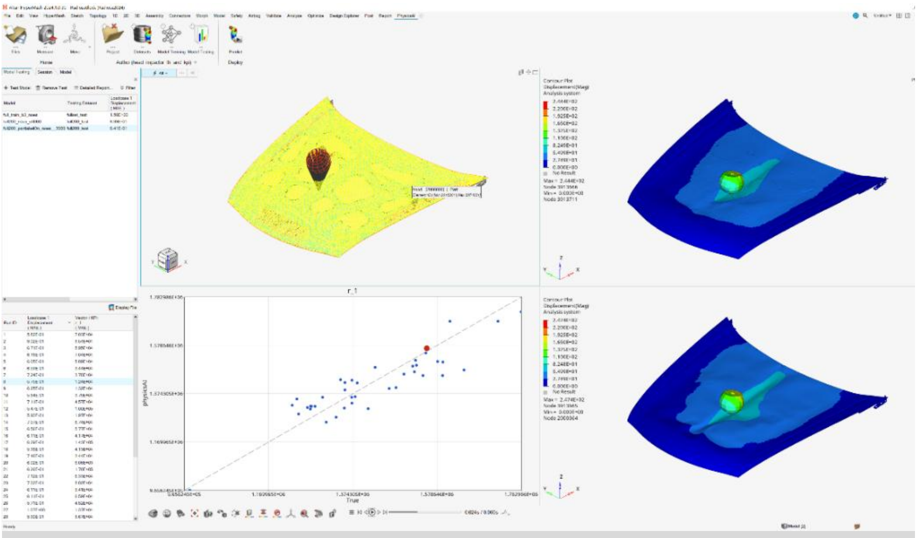

Figure 6 – The PhysicsAI prediction user experience in Altair HyperMesh. The comparison shows excellent agreement between true (top right) and predicted (bottom right) displacement field contours. The graph (lower left) shows a scatter of true vs. predicted maximum acceleration. The high correlation is indicated by the proximity of all data points to the diagonal.

Comparing the Effect of HPC Hardware on the AI Training Process

While PhysicsAI models can be trained on any machine, the training process will be significantly quicker when trained on a GPU as opposed to a CPU. The types of CPUs and GPUs engineers use for this process also have an impact. To compare, Table 1 lists the training times for several configurations. Simply moving from a workstation CPU to a workstation GPU reduced training times by a factor of three. A dedicated datacenter GPU, however, trained the model 18-times faster than a workstation CPU.

Table 1. A comparison of training times using PhysicsAI 2024.1 for the head impact dataset with 160 simulations in the training set and 3000 epochs. The workstation CPU is an Intel Core i7-12800H. The Workstation GPU is an NVIDIA RTX A2000. The datacenter GPU 1 is an NVIDIA A10. The datacenter GPU 2 is an NVIDIA A100-SXM

Configuration	Train Time (Day: Hour: Minute)	% reduction	Speed Up Factor
Workstation CPU	3:7:29	N/A	N/A
Workstation GPU	1:0:41	69%	3
Datacenter GPU 1	0:8:42	64%	9
Datacenter GPU 2	0:4:29	94%	18

This comparison illustrates the substantial time savings that come from GPU hardware. However, these savings must be balanced against scalability and budgetary concerns. Hence, Altair One's ability to give engineers access to scalable, on-demand, and affordable GPU infrastructure significantly improves the AI model training process.

Conclusion

For decades, the HyperWorks design and simulation platform has provided best-in-class simulation-driven design workflows for every industry. In that time, the Altair® HPCWorks® cloud and HPC platform have also facilitated companies to maximize the efficiency of compute resources. The Altair One cloud innovation gateway brings these platforms together to provide browser-based access to its tools, data, and compute resources. Available from within Altair One, the geometric deep learning of PhysicsAI combines data science and engineering to quickly provide insights and improve designs. Including these fast predictions into simulation-driven design workflows will enable engineers to solve the world's toughest problems.

To learn more sign up for Altair One, or request a free trial of Altair PhysicsAI.