

# Continuous Physics Reasoning

## Definition, Minimum Criteria, and the Role of Foundation Models for Physics

Hardik Kabaria, Sarah Osentoski

April 2026

Reference page: [Vinci Research](#)

Discussion / version history: [GitHub repository](#)

Suggested citation

Kabaria, H., Osentoski, S. Continuous Physics Reasoning: Definition, Minimum Criteria, and the Role of Foundation Models for Physics. Vinci Research, April 2026.

### Abstract

Physical products are increasingly constrained by thermal, mechanical, electrical, and manufacturing realities, yet much of industry still relies on intermittent, expert-mediated physics evaluation. As systems become more complex and tightly coupled, this limits not only product development, but also manufacturing readiness, operational efficiency, and lifecycle performance.

This paper proposes Continuous Physics Reasoning as a category of engineering capability in which physical reasoning becomes continuously available across the product lifecycle—from design, evaluation, and validation to manufacturing and operation—rather than being confined to discrete simulation checkpoints.

We define a foundation model for physics as the enabling system class for Continuous Physics Reasoning: a general-purpose system that reasons natively over physical structure with deterministic, solver-grade, out-of-the-box generality at manufacturing resolution.

We distinguish necessary from sufficient conditions, propose a minimum qualifying standard, and identify disqualifying patterns such as task-specific surrogates, per-domain model forks, nondeterministic execution, and workflow automation without general-purpose physical reasoning.

The goal is to provide a classification framework for evaluating systems that claim to enable Continuous Physics Reasoning.

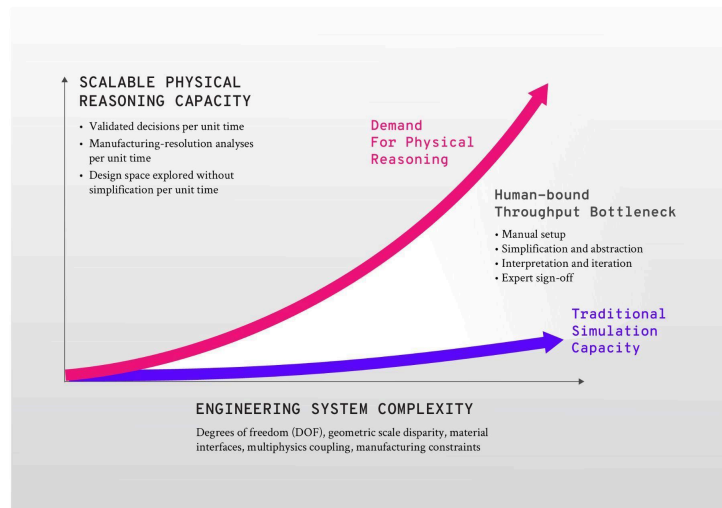
# 1. Executive Summary

The way physical products are designed, tested, and built is approaching a structural limit. Across industries—from electronics and automotive to medical devices, energy systems, and aerospace—engineers are being asked to reason about a larger number of systems that are more complex, tightly coupled, and constrained by manufacturing realities than ever before. While simulation, automation, and compute have improved individual steps in the engineering workflow, physical reasoning remains episodic, labor-intensive, and difficult to scale.

At the same time, advances in artificial intelligence have demonstrated that foundation models can serve as general-purpose reasoning layers for complex domains. In language and perception, foundation models reshaped work by enabling continuous, interactive reasoning rather than task-by-task execution. This

document argues that an analogous shift is now possible for the physical world: a new category of engineering capability, Continuous Physics Reasoning, in which physical reasoning becomes continuously available across design, evaluation, and validation rather than being confined to intermittent simulation checkpoints. Realizing this category requires an enabling system class: a foundation model for physics. This document defines a foundation model for physics as the enabling system class for Continuous Physics Reasoning: a

general-purpose intelligence layer that reasons natively over physical structure—geometry, materials, and boundary conditions—and operates deterministically at manufacturing resolution. Unlike traditional simulation tools and most physics AI approaches, a true foundation model for physics must generalize across physical systems—whether human-engineered or naturally occurring—governed by the same underlying laws of physics— without requiring customer or user-specific training, tuning, or adaptation. This “out-of-the-box physical understanding” is the defining requirement that separates foundation models from task-scoped, data-dependent, or configuration-heavy systems, and makes Continuous Physics Reasoning possible. [1]

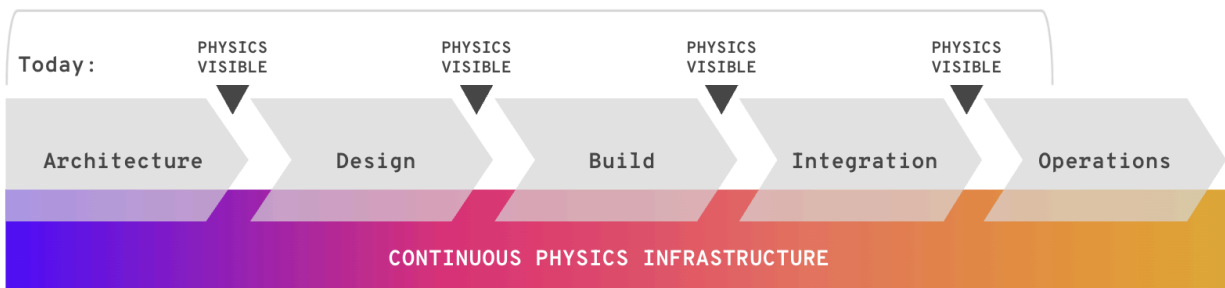


Demand is exploding because modern products concentrate more physics across more scales, under tighter manufacturing and reliability constraints — while the cost of being wrong keeps rising.

The purpose of this document is to:

- provide a precise definition of the category of Continuous Physics Reasoning
- define the enabling role of foundation models for physics
- establish a minimum qualifying bar
- clarify what does and does not qualify
- outline the implications of adopting Continuous Physics Reasoning for how physical products are conceived, evaluated, and brought into the world

The goal is not to promote a specific implementation or certify any current system, but to provide a classification framework that enables consistent evaluation and clearer discussion across industry, research, and analyst communities.



## 2. Why a New Category Is Required

Most engineering workflows for physical products are built around a structural assumption: that physics is evaluated intermittently, rather than continuously, throughout the development process. Physical reasoning is treated as a series of discrete checkpoints—invoked at specific stages to validate designs—rather than as an always-on capability that guides exploration and decision-making.

This paradigm worked when products were simpler, physical coupling was limited, and expert-driven analysis could scale with complexity. In practice, understanding the physical behavior of a part or system was performed sporadically, by a relatively small population of highly specialized engineers, at discrete points in the development cycle.

Today's physical systems are fundamentally different. They involve high-dimensional geometry, heterogeneous materials, tightly coupled multiphysics effects, and manufacturing constraints that must be reasoned about simultaneously. In this regime, the primary bottleneck is no longer solver

capability or compute availability, but the scarcity and throughput of human-mediated physical reasoning—the effort required to translate design intent into analyzable physics and back into actionable decisions.

As a result, physics has remained a periodic checkpoint rather than an always-on capability, limiting not just speed, but the breadth of designs that can be explored, evaluated, and ultimately built.

A range of AI-based approaches has emerged to address parts of this problem—workflow automation, surrogate models trained on simulation or experimental data, and physics-informed machine learning. While valuable in specific contexts, these approaches often inherit the same structural limitations: they remain configuration-heavy, require training or tuning on representative data, or trade determinism and validation for speed. Most importantly, they tend to treat AI as an adjunct to existing workflows rather than as a persistent reasoning layer for physical reality itself. [2-7]

What is missing is a category of engineering capability in which physical reasoning becomes continuously available, enabled by a model class that can act as a general-purpose physical reasoning layer, analogous to how foundation models operate in other domains. Such a system would support continuous interaction with physical reality—enabling rapid exploration of design spaces, evaluation of tradeoffs, and validation of decisions without repeated retraining, reconfiguration, or manual intervention for each new problem.

| KEY DISTINCTION FOR YOUR CATEGORY WORK                                      |   |
|---|---|
| TRADITIONAL SOLVER  | PHYSICS FOUNDATION MODEL  |
| <ul style="list-style-type: none"><li>• Solves physics accurately</li></ul> | <ul style="list-style-type: none"><li>• Reasons over physics continuously</li></ul> |
| <ul style="list-style-type: none"><li>• Requires setup per case</li></ul>   | <ul style="list-style-type: none"><li>• Out-of-the-box operation</li></ul>          |
| <ul style="list-style-type: none"><li>• Human-bound workflows</li></ul>     | <ul style="list-style-type: none"><li>• Machine-scale reasoning</li></ul>           |
| <ul style="list-style-type: none"><li>• Episodic runs</li></ul>             | <ul style="list-style-type: none"><li>• Always on</li></ul>                         |
| <ul style="list-style-type: none"><li>• Tool</li></ul>                      | <ul style="list-style-type: none"><li>• Intelligence layer</li></ul>                |

Existing categories do not adequately describe or constrain this kind of system, which is why a new definition—and a minimum qualifying bar—are required.

### **3. Definition of Continuous Physics Reasoning and the Role of a Foundation Model for Physics**

Continuous Physics Reasoning is a category of engineering capability in which physical behavior can be evaluated continuously across the product lifecycle. A foundation model for physics is the enabling system class for Continuous Physics Reasoning, providing deterministic, solver-grade reasoning over full-fidelity physical systems at manufacturing resolution without task-specific or customer-specific training.

At its core, a foundation model for physics functions as an always-on intelligence layer for physical reasoning. It operates directly on representations of physical structure (geometry, materials, boundary conditions) and produces outputs that are reproducible and verifiable for robust decision-making in production engineering, scientific, and exploratory contexts. Its defining role is not simply to accelerate known analyses, but to enable engineers to reason over a broader set of physical possibilities than would be practical through human-bounded workflows alone. [7, 8]

A foundation model for physics is defined by deterministic solver-grade reasoning, scalable physics-native representation, and out-of-the-box generalization as a single global model. Section 5 formalizes these requirements as a “minimum qualifying standard.”

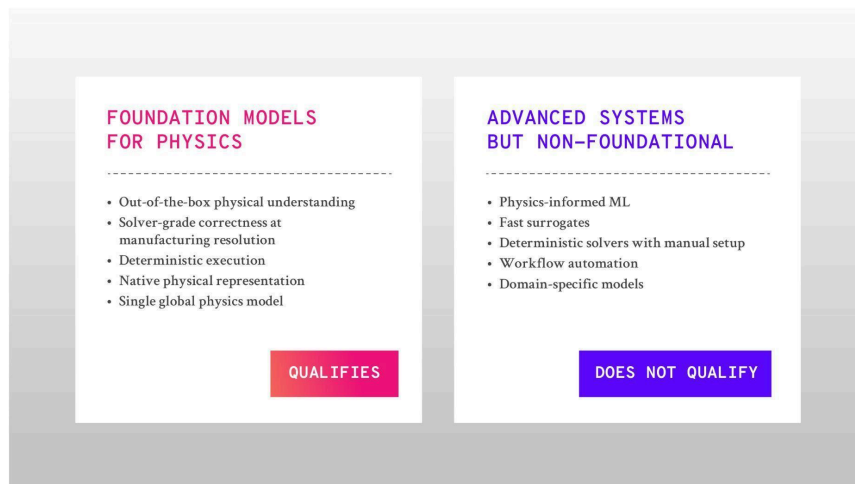
This definition distinguishes foundation models for physics from adjacent system classes. Traditional physics simulation tools can be deterministic and solver-accurate but are limited by human-bounded workflows. AI-assisted engineering tools can automate portions of the workflow but often remain configuration-heavy and domain-specific. Neural surrogates can be extremely fast, but they are often trained for specific operating regimes and may not generalize outside their training envelope and may not meet determinism or validation requirements. The next section clarifies why several commonly cited “physics AI” properties are necessary but not sufficient, and why a stricter minimum bar is required.

## 4. Necessary vs. Sufficient Conditions

As AI enters physical engineering, certain capabilities are frequently cited as indicators of progress—deterministic outputs, physics-informed learning, and validation against first-principles solvers. These properties are meaningful and increasingly expected in serious engineering contexts. However, they are necessary but not sufficient to qualify a system as a foundation model for physics. [2, 3, 5]

Determinism ensures repeatability but does not ensure generality. Physics-informed learning can embed constraints but does not guarantee that a model can reason beyond its training distribution. Validation frameworks can establish correctness for specific tasks while still requiring extensive tuning, reconfiguration, or domain-specific adaptation. Systems that satisfy these conditions may improve existing workflows, but they remain fundamentally task-scoped, training-dependent, or human-bounded.

The distinction matters because foundation models in other domains are not defined merely by scale or performance, but by their ability to serve as general-purpose reasoning layers across tasks without adaptation. The same principle must apply to systems that claim to enable Continuous Physics Reasoning. A foundation model for physics must meet a higher bar: it must support general-purpose physical reasoning without requiring retraining, tuning, or per-problem configuration to operate on new physical systems.



Section 5 formalizes this higher bar as a minimum qualifying standard.

## 5. Minimum Bar: Qualifying Criteria for Enabling Continuous Physics Reasoning

A system qualifies as enabling Continuous Physics Reasoning in the strong sense only if it satisfies all six criteria below. In this document, a foundation model for physics is the enabling system class that meets this standard. These criteria are intentionally strict. Failure to meet any single criterion disqualifies a system from reaching the minimum bar for enabling Continuous Physics Reasoning, regardless of benchmark performance or domain success.

### 5.1 Physics-Native Representation and Reasoning

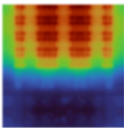
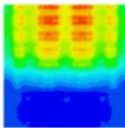
The model reasons directly over physical structure including geometry, materials, the physical forces, fluxes, and operating conditions that act on a system rather than abstracted features, reduced-order surrogates, or task-specific encodings. Physical understanding must be intrinsic to the model rather than inferred indirectly through abstracted features or task-specific encodings.

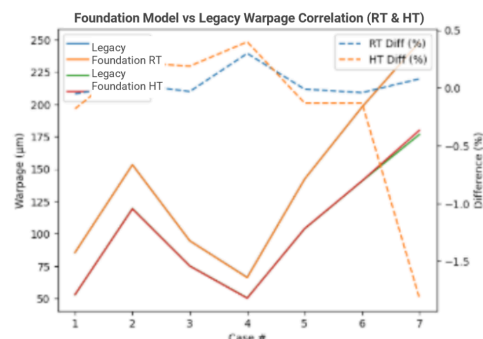
### 5.2 Deterministic Physical Reasoning

Identical physical inputs produce identical outputs—deterministic for classical systems or probabilistic for stochastic processes like nanoscale manufacturing—enabling reproducibility via statistical measures, uncertainty quantification, regression for optimization, and empirical validation. Reproducibility must be intrinsic—not achieved through averaging, post-processing, or statistical stabilization.

### 5.3 Solver-Grade Correctness at Manufacturing Resolution

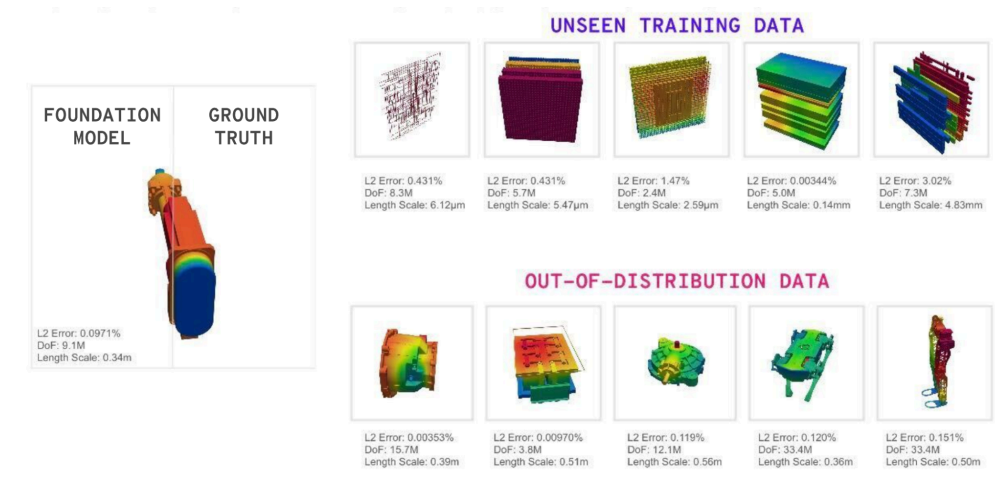
Outputs must be continuously verifiable against first-principles solvers or equivalent physical ground truth at manufacturing resolution. The model natively respects fundamental physical invariants (e.g., conservation of mass, energy, momentum; thermodynamic consistency; symmetries) across all regimes, including extrapolation. Violations are prevented by design or reliably detected and flagged with explanations. Approximate correctness or empirical accuracy alone is insufficient; validation must support production-grade trust.

|              | Foundation Model  | Commercial FEA Solver   |
|--------------|---|---|
|              |  |  |
| Max          | 100.32°C  | 101.11°C  |
| Avg          | 56.49°C   | 56.50°C   |
| Min          | 26.18°C   | 25.93°C   |
| Meshing      | Automated   | Manual, 30 Mins   |
| Compute time | 20 SECONDS  | 2.5 HOURS   |



## 5.4 Out-of-the-Box Physical Understanding

The model must generalize across new geometries, materials (including heterogeneous, anisotropic, and coupled behaviors), boundary conditions, and product types without customer-specific training, fine-tuning, or adaptation. Inference must not depend on prior exposure to similar designs.



## 5.5 Universality and a Single Global Physics Model

To enable Continuous Physics Reasoning, a foundation model for physics must operate as a single, unified reasoning model within its declared physical scope.

Unlike language or software domains—where symbols, abstractions, and rules are human-defined—physics is not constructed or negotiated. Physical systems interact according to universal laws that apply automatically and consistently at the appropriate scale. Whether governing heat transfer, electromagnetic fields, structural mechanics, or conservation of momentum, these laws exist independently of human interpretation and are expressed continuously through matter itself.

Because physics is universal, a foundation model for physics cannot be fragmented by customer, domain, or application. Per-customer models, per-domain forks, custom code paths, or forward-deployed engineering for training or tuning fundamentally contradict the notion of foundational physical understanding. Such approaches imply that physical behavior must be relearned or reinterpreted for each context, rather than reasoned about as a consequence of shared underlying laws.

A single global model is therefore not an architectural preference, but a defining requirement. It reflects the reality that physical parts and systems—whether electronic, mechanical, biological, or structural—interact through the same governing principles, even as scale, materials, and operating conditions vary. Generalization must arise from intrinsic understanding of those principles, not from empirical specialization or model proliferation.

This criterion distinguishes foundation models for physics from collections of specialized tools or models optimized for narrow regimes. A system that requires customization to operate across physical domains is not reasoning over physics itself; it is assembling approximations. Only a single global model can serve as a true, general-purpose reasoning layer for the physical world.

## 5.6 Scalable Efficiency

The model operates at high resolutions and scales (e.g., from nano to macroscopic) with low computational overhead, enabling real-time inference for complex systems without sacrificing accuracy.

Together, the criteria in Section 5 define a minimum and intentionally disqualifying bar for foundation status. They are orthogonal to performance metrics, workflow tooling, and interface choices. Such properties may emerge as consequences of meeting this bar, but they do not define it. Section 6 clarifies which classes of approaches do not qualify under this standard.

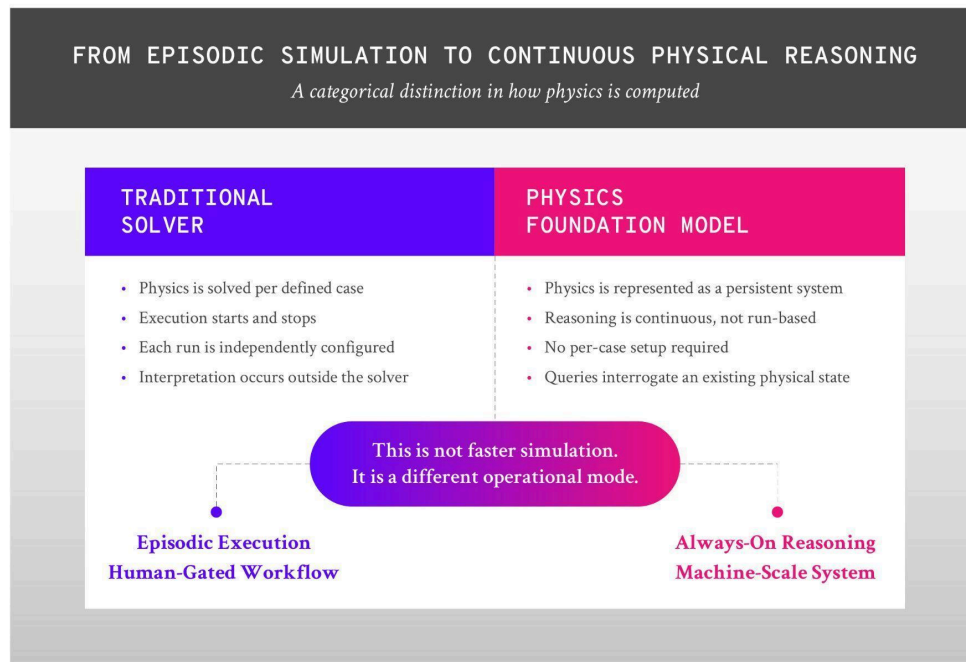
## 6. What Does Not Qualify as Continuous Physics Reasoning

Because “physics AI” is a broad label, it is useful to clarify disqualifying patterns. Systems that fall into any of the categories below may improve physics-driven workflows, but they do not qualify as enabling Continuous Physics Reasoning under the minimum bar defined in Section 5.

- **Training-dependent systems** that *require* customer-specific training, tuning, or adaptation to operate on new designs or domains.
- **Per-domain or per-customer variants** that rely on model forks, custom code paths, or fragmented architectures.
- **Surrogate-only or reduced-order approaches** that approximate solver behavior within a bounded design space without solver-grade correctness and out-of-the-box generalization.

- **Nondeterministic model execution** that produces outputs which are not exactly reproducible for identical physical inputs and conditions, excluding systems whose underlying physical formulation is explicitly probabilistic but still solver-grade.
- **Unphysical or invariant-violating models** that systematically break fundamental physical laws (e.g., conservation of mass, energy, momentum, or charge) outside of controlled approximations. A foundation model must either enforce such invariants by design or reliably detect and flag violations.
- **Workflow automation layers** that improve simulation pipelines but do not provide general-purpose physical reasoning.
- **Purely empirical black-box models** trained on simulation or experimental data without any explicit incorporation of physical constraints, symbolic equations, or solver supervision during training or inference.

These approaches may deliver meaningful improvements over traditional workflows, but they do not meet the minimum bar required for enabling Continuous Physics Reasoning.



Section 7 describes what changes when the minimum bar is met.

## 7. Behavioral and Workflow Implications

When Continuous Physics Reasoning is achieved, its impact transcends incremental gains—reshaping engineering by expanding who can ask physical questions, and transforming how those questions are posed, explored, validated, and executed.

- **From episodic physics evaluation to Continuous Physics Reasoning:** physics becomes continuously available throughout design and development rather than a periodic checkpoint.
- **From manual setup to automated exploration:** the cost of asking physical questions drops, shifting effort from configuration to exploration and tradeoff analysis.
- **From post-hoc validation to physics-authoritative validation by construction:** determinism and physics-grade correctness make validation intrinsic to the reasoning process—closer to just-in-time compilation of a design than after-the-fact verification—enabling earlier confidence and clearer comparison.
- **From disconnected tools to shared physical context:** diverse stakeholders can ask different questions against a consistent underlying representation of physical reality rather than through isolated toolchains.
- **From human-bounded iteration to machine-scale reasoning:** effective speed, multi-resolution analysis, and flexible outputs emerge as consequences of continuous physical reasoning at scale.
- **From precedent-based design to possibility-expanding discovery:** engineers are no longer constrained to configurations shaped primarily by historical patterns, existing templates, or human intuition. Broad out-of-the-box generalization enables exploration of unconventional physical regimes and novel design frontiers, unlocking solutions that would be unlikely to emerge through incremental iteration alone.

These shifts mirror the pattern observed when foundation models were introduced in other domains: the transformation comes from changing the order of operations, not from replacing one tool with another.

## 8. Domains, Generalization, and Initial Proving Grounds

While Continuous Physics Reasoning is a domain-agnostic category, credibility for the category will be established first in domains where physical reasoning is most demanding and where correctness and trust are contractual requirements rather than aspirational goals. Semiconductors and advanced electronic systems represent one of the most stringent proving grounds: they combine nanometer-scale features with centimeter-scale assemblies, dense material heterogeneity, and tightly coupled thermal, mechanical, and electrical effects. [7, 8]

Accurate predictions in this regime require manufacturing-resolution fidelity across extreme scale ratios while preserving solver-grade correctness—not only to achieve technical insight, but to satisfy the contractual, qualification, and sign-off obligations that govern real production hardware. In addition, these domains impose some of the strictest requirements around data secrecy and IP protection, where design data cannot be shared, abstracted, or externalized without risk. Any system that succeeds here must therefore operate with production-grade correctness and within the security, privacy, and trust constraints of world-class hardware organizations.

In complex domains like semiconductor, aerospace, and material engineering, human-bounded workflows are particularly limiting. Manual tasks such as geometry meshing, boundary setup and solver configuration constrain throughput and stifle broad exploration of design spaces. Data-driven surrogate approaches exacerbate this by demanding human intervention to curate representative datasets, orchestrate training regimes, and iteratively retrain as geometries, materials and processes evolve. As a result, physical reasoning remains bounded by human throughput, and the gap between what is physically relevant and what is practically analyzable continues to widen.

Demonstrating Continuous Physics Reasoning in such environments offers compelling evidence of true foundational capability. A system that can reliably reason over complex, multi-scale physical structures—spanning nano defects in wafers to macroscopic fluid structure interactions — without retraining, tuning, or per-design adaptation has inherently internalized structural physical understanding rather than pattern-specific behavior.

Success in semiconductors and electronics is therefore significant not because these domains are the endpoint, but because they are a rigorous benchmark. The same foundational properties required here apply broadly across physical domains, including automotive systems, medical devices, energy infrastructure, aerospace, and industrial equipment—within the model’s defined scope.

## 9. Implications for Industry, Research, and Evaluation

Continuous Physics Reasoning changes how engineering systems should be evaluated. Rather than assessing tools solely on benchmark performance or domain features, organizations should ask whether a system meets the minimum bar to credibly enable Continuous Physics Reasoning and whether its generality compounds over time. In this framing, numerical solvers serve as instruments of verification, while physical authority resides in the model's invariant, deterministic reasoning over governing laws.

For researchers, foundation models for physics elevate questions of representation, validation, and generalization under strict determinism and solver-grade correctness. They also sharpen the distinction between empirical interpolation and structural reasoning—an area that will shape future work at the intersection of AI, numerical methods, and physical sciences. [8]

For analysts and evaluators, a precise definition reduces category dilution. Without a shared standard, disparate approaches are grouped under broad labels (“physics AI,” “AI-assisted simulation”), obscuring meaningful differences. The criteria in this document provide a defensible framework for classification and comparison.

Finally, governance matters. Determinism, solver-grade validation, and explicit scope boundaries are prerequisites for responsible deployment in high-stakes contexts. As adoption grows, these factors will increasingly influence certification, auditability, and trust.

## 10. Conclusion: Toward Continuous Physics Reasoning

Continuous Physics Reasoning is not defined by speed, scale, or tooling alone. It is made possible by a foundation model for physics defined by its ability to function as a general-purpose reasoning layer for physical reality: reasoning natively over physical structure, running on new designs out-of-the-box—without retraining, tuning, or calibration loops, and doing so as a single unified model. Without a strict definition, the term risks dilution. With a disqualifying minimum bar, researchers, analysts, and practitioners can evaluate claims consistently and distinguish between incremental improvements and a genuine category shift.

When the minimum bar is met, the consequences extend beyond simulation. Physics becomes continuously accessible, exploration becomes the dominant mode of interaction, and validation becomes intrinsic rather than retrospective. Over time, this shift can reshape how physical products are conceived, evaluated, and brought into the world.

The intent of this document is not to predict the pace of that transition, but to provide a clear framework for recognizing it—and for reasoning about what follows—when foundation models for physics begin to emerge across industries.

This definition and framework for Continuous Physics Reasoning is intended to establish a clear category standard for evaluating systems that claim to enable it.

## Acknowledgments

This paper benefited from review and critique from external researchers, practitioners, and analysts. Their participation helped sharpen the framework, but does not imply agreement with every claim or endorsement of the paper as a whole.

## References

- [1] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the Opportunities and Risks of Foundation Models. arXiv preprint arXiv:2108.07258, 2021.
- [2] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [3] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.
- [4] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaijun Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations. arXiv preprint arXiv:2010.08895, 2020.
- [5] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations. arXiv preprint arXiv:1711.10561, 2017.
- [6] Ehsan Haghighat, Maziar Raissi, Alberto Moure, Hector Gomez, and Ruben Juanes. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 379:113741, 2021.
- [7] G. Refai-Ahmed, M. M. Islam, M. Shahsavan, H. Do, H. Kabaria, J. L. Davenport, J. G. Kocheemoolayil, N. Harrington, S. Osentoski, and S. D. Beer Mohideen, “Towards AI-Assisted Design of Thermal Management Strategies,” in 2024 IEEE 26th Electronics Packaging Technology Conference (EPTC), 2024.
- [8] H. Kabaria, S. D. Beer Mohideen, J. L. Davenport, B. Cherukuri, J. Kocheemoolayil, N. Patil, S. Osentoski, W. Stark, and A. J. Lew, “Thermal Sensitivity Analysis of 3D IC Face-to-Back Stacking Using Foundation Models for Physics,” 2025.