Mechanics in Nano-Engineering 1(1) (2026) 1-18

Ministry of Science, Research and Technology
University of Bonab

Mechanics in Nano-Engineering

journal homepage: https://mneng.ubonab.ac.ir/

Mechanics in NANO-ENGINEERING

Machine Learning–Accelerated Nanomechanical Design: A Design-to-Venture Readiness Framework and Techno-Economic Model for Nano-Architected Material Ventures

Shiva Mehdizadeh Aghdam

Department of Management and Accounting, Technical and Vocational University (TVU), Tehran, Iran
Corresponding author: Sh.Mehdizadeh@gmail.com

ARTICLE INFO

Article history:
Received 9 20 2026
Received in revised form 10 15 2026
Accepted 11 6 2026

Keywords:
Nano-architected materials
Machine learning
Mechanical metamaterials
Active learning
Techno-economic analysis
Technology readiness
Learning curve

ABSTRACT

Nano-architected materials, including nanolattices, truss- and shell-based metamaterials and spinodoid structures, can combine stiffness, strength, low density and energy absorption in ways that are difficult to achieve with monolithic materials. Machine learning can search these design spaces more efficiently than trial and error, but computational speed alone does not establish commercial viability. This structured narrative review asks when machine-learning-assisted nanomechanical design can reduce the cost and time of reaching a qualified product, and what evidence a venture needs before that reduction is credible. We examine forward surrogates, inverse design, physics-informed learning, machine-learned interatomic potentials and active learning, and the barriers between a simulated design and a qualified product. We propose a six-gate Design-to-Venture Readiness framework with indicative pass/fail criteria linked to technology readiness levels, and a techno-economic model that values fewer physical experiments, earlier market entry and the production learning curve. The model excludes performance premiums and willingness to pay, so it gives a conservative cost-side value. With hypothetical inputs, the break-even speed-up rises from 1.7 to 3.4 times when 40% of avoided experiments must be repeated, and production volume, not experimental savings, is the binding constraint. All parameters are illustrative and require calibration with venture data.

1. Introduction

1.1. Nano-architected materials and their mechanics

Over the past decade, micro- and nanoscale architecture has emerged as a route to lightweight mechanical performance. Ceramic nanolattices can recover after compressions that would fracture bulk ceramics [1]; glassy-carbon lattices approach the theoretical strength of their constituent material [2]; plate-nanolattices approach theoretical stiffness and strength limits at a given density [3]; and nanoarchitected carbon can dissipate energy under supersonic microparticle impact [4]. Together, these studies establish nanolattices as an important class of mechanical metamaterials [5,6].

In cellular solids, effective stiffness and strength scale with relative density and topology, and stretching-dominated lattices outperform bending-dominated ones at low density [7,8]. At the nanoscale, further effects enter. Strength and plasticity depend on specimen size [9], surfaces dominate the response, and fabrication defects can pull measured properties well below ideal predictions. Early demonstrations such as ultralight metallic microlattices [10] and ultrastiff metamaterials [11] showed what geometry can do, but they came from processes of limited throughput. The gap between idealized and measured behavior, usually expressed as a knockdown factor, is therefore central to any commercial claim [3,5].

1.2. Machine learning for nanomechanical design

Machine learning has become an established tool in materials research [12]. In mechanical metamaterials, generative models have been used to invert structure–property relationships and design nonlinear responses [13], and recent reviews survey both forward prediction and inverse-design workflows [14–17]. Four method families are particularly relevant to nanomechanical design.

Forward surrogates. Regressors trained on finite-element or atomistic data predict effective properties from geometry in milliseconds. Machine learning has been coupled with simulation and additive manufacturing in a closed loop for bioinspired hierarchical composites [18], and uncertainty-aware, data-driven frameworks propagate microstructural variability to component-level response [19].

Inverse design. Rather than searching geometry space, inverse models map target properties to structures. Deep learning has inverted the structure–property map of truss metamaterials [20], and inverse-designed spinodoid metamaterials offer wide tunability of anisotropic stiffness [21]. These approaches complement gradient-based topology optimization [22] by handling larger design spaces and non-differentiable objectives.

Physics-informed learning and machine-learned interatomic potentials. Embedding governing equations in the loss function can lower data requirements and may improve extrapolation [23,24], which matters when every data point is a costly experiment. Machine-learned interatomic potentials can approach density-functional-theory accuracy in atomistic simulation at far lower cost than ab initio calculation [25,26]. They enter the pipeline of Section 2.2 at the data-generation stage and mainly at the nanoscale. For struts or shells with characteristic dimensions of tens to a few hundred nanometres, atomistic simulation can supply size-dependent elastic constants, surface effects and defect energetics to continuum finite-element models. For microscale lattices (struts of roughly a micrometre or more), which account for most of the experimentally validated architectures cited here, continuum models with measured constitutive properties are the main data source and interatomic potentials play at most a supporting role. We use “nano-architected” in its common, loose sense; the distinction matters because it determines which simulator supplies training data and how large the size-effect knockdown is likely to be.

Active learning and autonomous experimentation. Adaptive sampling uses model uncertainty to choose the next most informative simulation or test [27]. A Bayesian autonomous experimenter has been demonstrated for mechanical design, iteratively fabricating and testing structures to improve energy absorption [28], and Bayesian optimization has been used to exploit, rather than merely avoid, fabrication defects in architected materials [29].

1.3. Commercialization frameworks and the research gap

For a start-up, faster design cycles may make it possible to explore more candidate architectures with a small team. Yet a simulated geometry is not a product: customers need manufactured components that meet performance requirements at an acceptable cost. The transition from invention to a viable business remains difficult for deep-technology ventures [30]. Technology readiness levels [31], stage-gate reviews [32], business-model design [33], appropriability theory [34] and customer-validated learning [35,36] describe how ventures develop and capture value, but they leave a gap. Technology readiness levels grade maturity without saying what computational evidence is sufficient at each level. Generic stage-gate reviews ask whether a business case is sound without asking whether a predictive model has been validated on physical tests, how much of its promised saving survives that validation, or how a saving in experiments translates into venture value.

The Design-to-Venture Readiness framework proposed here differs from these approaches in three ways. It specifies model-validity evidence, namely calibrated uncertainty on held-out physical tests and an estimate of the discrepancy fraction φ. It ties each gate to measurable acceptance criteria and to technology readiness level. And it links the gates to a quantitative model of cost-side value.

Research gap and contribution. Existing technology-readiness and stage-gate frameworks assess maturity and business readiness, but do not specify how held-out physical validation, model discrepancy, manufacturing learning and venture runway should be linked in one decision process for nano-architected materials. The Design-to-Venture Readiness framework operationalizes that link through six evidence-based gates and couples the gates to a cost-side model; it is a decision aid, not a substitute for market validation.

1.4. Scope and objectives

The economic analysis in this article values only the physical experiments avoided and the margin earned by reaching the market earlier, together with the production economics needed to reach a target unit cost. It does not value performance premiums, new-market creation, customer willingness to pay, competitor response or pricing strategy. Conclusions about viability are therefore conditional on price being an exogenous input, and the model gives a conservative cost-side lower bound on the value of machine learning, not a full business-case valuation.

The article addresses three questions: (1) Which machine-learning approaches can accelerate nanomechanical design while remaining useful after experimental validation? (2) Which barriers separate computational gains from commercial gains? (3) How can a venture estimate whether, and where, investment in machine learning is justified on cost-side grounds? It contributes a structured narrative review (Section 2.1), a techno-economic model (Section 2.3) and the Design-to-Venture Readiness framework (Section 2.4), and evaluates them with illustrative parameters (Section 3).

2. Materials and Methods

2.1. Review design and search protocol

This study is a structured narrative review combined with conceptual framework development. It is not a systematic review: sources were selected purposively, not through an exhaustive database protocol, so the evidence base may be incomplete (Section 3.8). The selection steps are described with reference to the PRISMA 2020 reporting items that apply [37].

Sources were identified through targeted searches in Google Scholar and publisher platforms, supplemented by backward and forward citation tracking from recent reviews [14–17]. Architected-material terms (nanolattice, architected material, mechanical metamaterial, spinodoid, plate lattice) were combined with machine-learning terms (machine learning, neural network, inverse design, Bayesian optimization, active learning, surrogate); commercialization terms (technology readiness, stage-gate, learning curve, experience curve, techno-economic) were searched separately. The time window was 2011 to 2026, starting from the first ultralight microlattice demonstration [10]. The last search was made in October 2026. Standards and books were added by hand.

Included were peer-reviewed articles, reviews, standards and books that (a) apply machine learning or data-driven optimization to the mechanical design of architected, cellular or truss/shell/plate metamaterials; or (b) report fabricated and mechanically tested designs with a predicted-versus-measured comparison; or (c) treat technology readiness, stage-gate processes, learning curves or calibration methods relevant to the model. Excluded were optical, acoustic and electromagnetic metamaterials, non-English papers and preprints not later published. Simulation-only studies were retained to describe methods but were excluded from the evidence table. Screening proceeded from titles and abstracts to full texts by the author alone, so no independent second screening was possible. For each fabricated-and-tested study we extracted the system and scale, what was physically tested, the reported benefit (for example, experiments saved) and any predicted-versus-measured gap, knockdown factor or estimate of φ.

2.2. A machine-learning-accelerated design pipeline

The venture-oriented pipeline has six technical components. The six readiness gates of Section 2.4 are venture-level checkpoints rather than pipeline stages; Table 1 and Fig. 1 map one onto the other.

Data generation and curation. Simulation sets (finite element, molecular dynamics) and experimental sets (nanoindentation, micro-compression, tomography of printed parts) must be stored with complete metadata on geometry and process parameters. Data quality caps every downstream step: surrogates and interatomic potentials inherit the biases of their training data [25,26].

Surrogate modeling. The model class should match the data volume: Gaussian processes or ensembles for small sets, deep networks for large ones, with physics-informed constraints in the low-data regime [23,24].

Table 1. Mapping of technical pipeline components to readiness gates.

Pipeline component	Main output	Gate(s) informed
Data generation and curation	Versioned geometry–process–property dataset	D1 Data; refreshed after gate decisions
Surrogate modeling	Fast property predictions from simulation and experimental data	D2 Model validity
Uncertainty quantification	Calibrated prediction intervals; epistemic and aleatoric uncertainty	D2 Model validity
Simulation-to-reality calibration	Estimates of discrepancy δ(x), knockdown factor and φ	D2 Model validity
Manufacturing-aware inverse design	Producible candidates; estimates of c_1 and b	D3 Producibility; inputs to D4 Economics
Closed-loop validation	New physical data; re-tests of D2 and D3	Feeds D1; re-evaluates D2 and D3
Venture level, outside the technical pipeline	Cost, time-to-volume and runway (model of Section 2.3); third-party qualification; protection of designs, data and know-how	D4 Economics; D5 Qualification; D6 IP and data protection

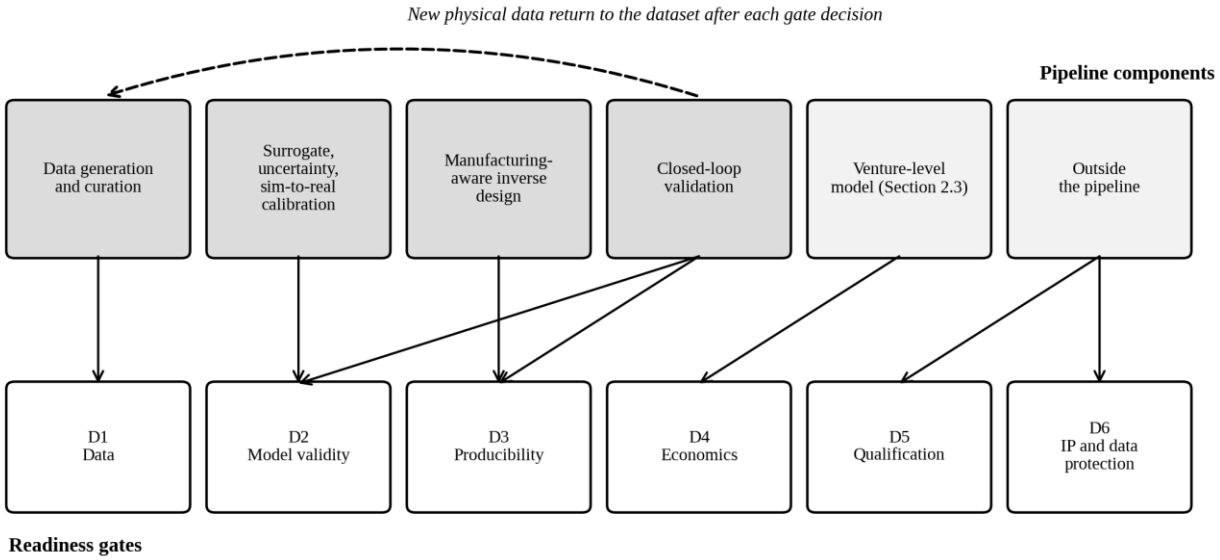


Figure 1. The machine-learning-accelerated design pipeline (top row) and the six Design-to-Venture Readiness gates (bottom row). Pilot test results return to the dataset after each gate decision; the held-out test set is locked during model selection.

Uncertainty quantification. A prediction without an uncertainty estimate cannot support an engineering decision. Practical options include deep ensembles [38], Monte Carlo dropout as approximate Bayesian inference [39] and Gaussian-process posteriors. Separating epistemic uncertainty (reducible with more data) from aleatoric scatter (fabrication variability) tells a venture whether to test more or to tighten its process.

Closing the simulation-to-reality gap. A standard formulation relates experiment to simulator as [40]

$$y_{\mathrm{exp}}(x) = \gamma\,\eta(x,\theta) + \delta(x) + \varepsilon \qquad \textbf{(1)}$$

where η is the simulator, θ its calibration parameters, γ a scaling factor, δ(x) a model-discrepancy term and ε measurement noise. Learning δ(x) from a modest number of physical tests can be a more

reliable use of a limited experimental budget than refining the simulator alone, because it targets the discrepancy the customer will actually observe. Whether this is also worth more commercially is an empirical question (Section 3.8).

Manufacturing-aware inverse design. Inverse models [13,20,21] should carry fabrication constraints in the objective: minimum feature size, support structures, printable orientation and post-processing shrinkage. Ignoring them risks a gap between predicted and manufactured performance, which is the mechanism behind the knockdown factors of Section 1.1 and the defect sensitivity reported in earlier work [3,29].

Closed-loop validation. Active learning [27] and autonomous experimentation [28] close the loop between design and test and direct physical experiments to where uncertainty is highest. The held-out test set is locked while a model is selected; once a gate decision is made, those measurements join the training data and a new held-out set is created for the next round. Additive-manufacturing terminology and process categories should follow established standards [41] so that results are reproducible across suppliers.

2.3. Techno-economic model of cost-side value

2.3.1. Variables

The model values the physical experiments avoided by machine-learning-assisted search, the margin earned by reaching the market earlier, and the production economics needed to reach a target unit cost. It does not value a performance premium, new-market creation, willingness to pay or competitor response, and the price p is an exogenous input. Table 2 defines every variable.

2.3.2. Value from fewer experiments and earlier market entry

A conventional campaign needs N_0 experiments. An ideal machine-learning campaign needs N_0/ρ and avoids $N_0(1 - 1/\rho)$. Let φ be the fraction of those avoided experiments that must be performed after all, because the model's predictions fail physical validation. The assisted campaign then performs $N_0/\rho + \varphi N_0(1 - 1/\rho)$ experiments, so the effective speed-up is

$$\rho_{\text{eff}} = \frac{N_0}{N_0/\rho + \varphi N_0(1 - 1/\rho)} = \frac{\rho}{1 + \varphi(\rho - 1)} \quad \textbf{(2)}$$

and the net cost-side value is

$$V_{\text{ML}} = N_0\, c_{\text{exp}}(1 - \varphi)\left(1 - \frac{1}{\rho}\right) + \delta R_m \Delta t - C_{\text{ML}} \quad \textbf{(3)}$$

Machine learning is justified on cost-side grounds when $V_{ML} > 0$. φ lies in [0, 1): $\varphi = 0$ means every predicted saving is realized and $\varphi \to 1$ means none is. From Eq. (2), $\rho_{\text{eff}} > 1$ if and only if $\varphi < 1$ (given $\rho > 1$), $\rho_{\text{eff}} \leq \rho$ always, and $\rho_{\text{eff}} \to 1$ as $\varphi \to 1$. Because the gain from machine learning is the number of experiments avoided, scaling the avoided fraction by $(1 - \varphi)$ is the definition consistent with the verbal meaning of φ. A formulation in which φ scales ρ itself, $\rho_{\text{eff}} = \rho(1 - \varphi)$, gives different results (Section 3.3) and can fall below 1, implying that machine learning slows the campaign, which is not what φ is meant to describe.

Table 2. Variables of the techno-economic model.

Symbol	Meaning	Unit
N_0	Physical experiments needed by a conventional campaign to reach a target design	Count
c_{exp}	Cost per physical experiment (specimen, fabrication, test, labour)	$ per experiment
ρ	Nominal speed-up: N_0 divided by the experiments an ideal machine-learning campaign would need	dimensionless, $\rho > 1$
φ	Fraction of the experiments predicted to be avoided that must nevertheless be performed because predictions fail physical validation (model discrepancy)	dimensionless, $0 \leq \varphi < 1$
ρ_{eff}	Effective speed-up after re-testing	Dimensionless
C_{ML}	One-off cost of data, model development, compute and specialist staff	$
Δt	Months saved in reaching market	Months
R_m	Monthly gross margin once on the market ($R_m = m_u \cdot K \cdot Q$ at steady state)	$ per month
δ	Discount-and-risk factor on the time-to-market benefit, $\delta = p_s(1 + r)^{-Tm}$, with p_s the probability of reaching market, r the monthly discount rate and T_m the months until the benefit accrues	dimensionless, $0 < \delta \leq 1$
V_{ML}	Net cost-side value of machine learning	$
θ	Break-even ratio, $\theta = (C_{ML} - \delta R_m \Delta t) / (N_0 c_{exp})$	Dimensionless
c_1	Cost of the first qualified part	$ per part
b	Progress ratio: unit cost is multiplied by b with every doubling of cumulative output	dimensionless, $0.5 < b < 1$
L	Learning rate, $L = 1 - b$	dimensionless (%)
α	Learning exponent, $\alpha = -\log_2 b$	Dimensionless
n	Cumulative count of qualified parts (a unit is one part of specified size and function that passes inspection)	Parts
c_n	Cost of the n-th qualified part	$ per part
p	Selling price	$ per part
m_u	Required unit margin	$ per part
n^*	Cumulative volume at which unit cost reaches $p - m_u$	Parts
n_0	Cumulative volume at which unit cost reaches p (zero unit margin)	Parts
C_{cum}	Cumulative production cost	$
L_{peak}	Peak net cash outlay on production if every part is sold at p	$
Q	Output of one production line	parts per month
K	Number of production lines	Count
T_{vol}	Time to produce n^* parts	Months
T_{dev}	Months of design, validation and machine-learning campaign before production	Months
R_w	Funding runway	Months
K_{avail}	Capital available	$
K_{req}	Capital required	$
F	Fixed investment (tooling, equipment, qualification)	$
n_h	Planning horizon	Parts
p_s	Probability of reaching market	Probability
r	Monthly discount rate	per month
T_m	Months until the market-entry benefit accrues	Months
Π	Profit from n parts, $\Pi(n; c_1, \alpha) = pn - c_1 n^{1-\alpha}/(1-\alpha) - F$	$

In the baseline, Δt is held constant on purpose: φ is applied only to experimental cost, so that the two channels can be seen separately. Re-testing would in practice also delay entry, and a sensitivity with $\Delta t_{eff} = (1 - \varphi)\Delta t$ is therefore reported. The factor δ discounts and risk-adjusts the market-entry benefit, which accrues later than the experimental savings and only if the venture survives;

experimental savings occur during the campaign and are left undiscounted. This timing convention is a limitation, and a venture with detailed cash-flow data should discount each cash flow separately. Setting $V_{ML} = 0$ gives the break-even speed-up

$$\rho_{\mathrm{BE}} = \frac{1}{1-\theta/(1-\varphi)}, \qquad \theta = \frac{C_{\mathrm{ML}} - \delta R_m \Delta t}{N_0\, c_{\mathrm{exp}}} \tag{4}$$

A finite break-even exists only when $\varphi < 1 - \theta$; beyond that, no nominal acceleration makes machine learning pay on cost-side grounds. The values $\varphi = 0$ to 0.6 are examined as scenarios because no published estimate of φ for nano-architected materials was found (Table 3). A venture should estimate φ from its own held-out physical tests (gate D2), for example as the share of screened candidates that fail validation and must be re-tested.

2.3.3. Unit economics along a learning curve

The cost of the n-th qualified part follows Wright's relationship [42],

$$c_n = c_1\, n^{-\alpha}, \qquad \alpha = -\log_2 b \tag{5}$$

so that $b = 0.85$ means cost falls 15% with every doubling of cumulative output (learning rate $L = 1 - b = 15\%$). For price p and required unit margin m_u, the unit cost reaches $p - m_u$ at the cumulative volume

$$n^* = \left(\frac{c_1}{p - m_u}\right)^{1/\alpha} \tag{6}$$

which is a cost threshold, not a cash break-even: it is where the marginal part becomes profitable enough, not where the venture has recovered its production spending. Treating n as continuous, the cumulative production cost is

$$C_{\mathrm{cum}}(n) \approx \frac{c_1 n^{1-\alpha}}{1-\alpha}, \qquad C_{\mathrm{cum}}(n^*) = \frac{(p - m_u)\, n^*}{1-\alpha} \tag{7}$$

If parts are sold at p from the first unit, the net outlay peaks where unit cost equals price, $n_0 = (c_1/p)^{1/\alpha}$, at

$$L_{\mathrm{peak}} = \frac{p\, n_0\, \alpha}{1-\alpha} \tag{8}$$

In practice little revenue arrives before qualification (gate D5), so $C_{cum}(n^*)$ is the more cautious figure.

Wright's relationship was derived for aircraft assembly [42]. Cost declines against cumulative output have since been found to be statistically regular across many technologies [43], and reported progress ratios in manufacturing cluster roughly between 0.7 and 0.9 [44]. We use it because it needs only two parameters, c_1 and b, that a venture can estimate from its own pilot builds. Its use for nano- and microscale fabrication has three limits. First, the evidence comes from high-volume manufacturing; we are not aware of experience-curve estimates for lattice fabrication by serial processes such as two-photon lithography. Second, learning in such processes may come in steps (parallel writing, faster scanning, automated post-processing, better yield) rather than smoothly, so Eq. (5) describes only an average trend. Third, cost can fall with cumulative output only if output rises, which depends on throughput.

The unit n is one qualified part, so yield losses raise c_1 and c_n instead of inflating volume. The value c_1 = \$1,000 is a placeholder for the first qualified part including materials, machine time, post-processing and inspection, to be replaced by pilot-build cost accounting. Learning rates of 10%, 15% and 20% are scenarios that bracket the commonly reported manufacturing range, with 10% as the cautious case; they are not estimates for this technology. The values p = \$150 and m_u = \$50 (target unit cost \$100) are hypothetical for a specialist component, and a combined throughput K·Q of 1,000 parts per month is assumed.

2.3.4. Combined decision equation and runway test

Manufacturing-aware design acts on c_1 and b; machine-learning-assisted search acts on N_0 and Δt. Both enter one value expression. Relative to a reference design and process (superscript ref) over a planning horizon of n_h parts, total cost-side value is

$$V_{\mathrm{tot}} = N_0 c_{\mathrm{exp}}(1-\varphi)\left(1-\frac{1}{\rho}\right) + \delta R_m \Delta t - C_{\mathrm{ML}} + \left[\Pi(n_h; c_1, \alpha) - \Pi\left(n_h; c_1^{\mathrm{ref}}, \alpha^{\mathrm{ref}}\right)\right] \quad \textbf{(9)}$$

The two channels are linked through $R_m = m_u KQ$: a lower unit cost raises the margin m_u, and throughput KQ sets how fast that margin accrues. The first three terms are V_{ML} of Eq. (3); the bracket is the value of the manufacturing channel. Value is not enough; the venture must also survive until n^* is reached. The runway test has a time condition and a cash condition:

$$T_{\mathrm{dev}} + \frac{n^*}{K\,Q} \le R_w \quad \textbf{(10)}$$

$$K_{\mathrm{req}} = C_{\mathrm{ML}} + \frac{N_0\, c_{\mathrm{exp}}}{\rho_{\mathrm{eff}}} + F + C_{\mathrm{cum}}(n^*) \le K_{\mathrm{avail}} \quad \textbf{(11)}$$

Eq. (11) assumes no revenue before n^*; $C_{cum}(n^*)$ is replaced by L_{peak} if sales begin at the first unit.

2.3.5. Uncertainty propagation

To show what propagating uncertainty means in practice, 200,000 draws were sampled from triangular distributions (minimum, mode, maximum): N_0 (100, 200, 400); c_{exp} (1,000, 2,000, 4,000); ρ (1.5, 3, 6); φ (0, 0.2, 0.5); Δt (0, 6, 9); R_m (10,000, 30,000, 60,000); C_{ML} (200,000, 350,000, 600,000); δ (0.3, 0.6, 1.0); b (0.80, 0.88, 0.95); c_1 (500, 1,000, 2,000); and KQ (200, 1,000, 3,000), with $p - m_u$ = \$100 and R_w = 24 months. These ranges are assumptions chosen for illustration, not estimates, and the inputs were sampled independently with a fixed random seed.

2.4. The Design-to-Venture Readiness framework

The framework defines six gates that should be passed before a venture escalates its commercial commitments. It complements technology readiness levels [31] and the stage-gate process [32]. D1 asks whether a curated, versioned dataset links geometry, process and measured properties. D2 asks whether predictions, with uncertainty, hold on unseen physical tests. D3 asks whether the designs can be made to tolerance with repeatability. D4 asks whether cost per part reaches the target price along a credible learning curve within the runway. D5 asks whether performance is independently verified. D6 asks whether the advantage is protectable through patents, defensible trade secrets and clear data rights. Intellectual-property protection (D6) is separated from qualification (D5) because

they need different evidence and fall at different times: priority filings should be made early, from about D2 onward, whereas third-party qualification is possible only once producible parts exist. D6 therefore runs in parallel to D3–D5 and must be complete before large customer commitments. Indicative numeric thresholds are given in Section 3.6.

3. Results and Discussion

3.1. Evidence from the literature

Table 3 summarizes what the fabricated-and-tested studies reviewed report. The reported benefit is usually expressed as fewer experiments or a better optimum. None of the studies extracted so far reports an estimate of the discrepancy fraction φ, and measured-versus-predicted gaps are reported unevenly. This is the main reason φ is treated as a scenario parameter throughout.

Table 3. Evidence on experimental savings and predicted-versus-measured gaps.

Study	System and scale	What was tested physically	Reported benefit; gap, knockdown or φ
Gongora et al. [28]	Parametric family of additively manufactured polymer structures; macro/mesoscale	Closed-loop printing and compression testing by an autonomous Bayesian system	About 60-fold fewer experiments than a grid search; gap not applicable (search driven by measured data); φ not reported
Snapp et al. [45]	Additively manufactured polymer cylindrical shells; 11-dimensional design space	More than 25,000 experiments selected by Bayesian optimization with human oversight	75.2% energy-absorbing efficiency, above the previously observed best of roughly 71–72%; φ not reported
Gu et al. [18]	3D-printed bioinspired hierarchical polymer composites; macroscale	Machine learning trained on a finite-element database of 100,000 composites; predicted designs printed and tested (three or more specimens per design)	Designs reported as tougher and stronger than training examples; physical validation reported; a comparable numerical prediction-error estimate was not reported in the source text
Vangelatos et al. [29]	Architected microlattices with designed defects (4×4×5 lattices)	One best intuitive design and eight randomly generated designs were fabricated and tested (nine in total); simulation was compared with measured critical buckling load	Optimum found with 250 cost-function evaluations in a design space of billions; normalized strain-energy density reported as 12,464 times that of the defect-free lattice; mean experiment–simulation difference in critical buckling load 7.4% (not an estimate of φ)
Bastek et al. [20]; Kumar et al. [21]	Truss and spinodoid metamaterials	Simulation-led inverse design; not used to estimate experimental prediction error	Not applicable
Crook et al. [3] (not a machine-learning study; reference point)	Carbon plate-nanolattices with wall thicknesses of about 150–260 nm	In situ compression, nano-computed tomography, Raman spectroscopy	Stiffness and strength reported as reaching the Hashin–Shtrikman and Suquet upper bounds; gains of up to 522% and 639% over beam-nanolattices; model error and φ not applicable

Simulation-led inverse-design studies, including truss and spinodoid work [20,21], are discussed in this review but are not used to estimate experimental prediction error unless fabricated-versus-

measured data are reported. Across the included studies, φ is not reported as a standardized quantity. The 7.4% difference reported by Vangelatos et al. [29] is a specific simulation-to-test discrepancy for critical buckling load; it should not be generalized to other designs or interpreted as φ.

3.2. Barriers between computational and commercial success

Experimental bottleneck. Nano- and microscale fabrication can have low throughput, and each validation specimen may be expensive. Machine learning can help select a smaller, more informative set of experiments, but it cannot eliminate physical validation.

Size effects and defects. Small-scale mechanical response can differ from bulk behavior [9], while fabrication defects may determine failure [3,29]. A surrogate trained only on idealized geometries may therefore overestimate performance in manufactured parts.

Data scarcity and ownership. Public datasets for the mechanics of nano-architected materials remain limited. Proprietary data may become a defensible asset [34], although datasets covering only a narrow design or process range may not generalize to new conditions.

Scale-up and unit cost. A high-performing architecture may still be commercially unattractive if production is slow or expensive. Unit cost often falls as cumulative output grows [42], but experience-curve evidence comes mostly from high-volume manufacturing [43,44]. A venture must therefore determine whether a learning curve can plausibly reach its target price, at an achievable throughput, before capital is exhausted (Section 3.4).

Certification and trust. Industrial customers need traceable test evidence rather than model predictions alone. Models can guide engineering decisions, but product claims require physical qualification.

3.3. Break-even speed-up and model discrepancy

With $N_0 = 200$, $c_{exp} = \$2{,}000$, $\Delta t = 6$ months, $R_m = \$30{,}000$ per month and $C_{ML} = \$350{,}000$ (all hypothetical), $\theta = 0.425$ at $\delta = 1$. The break-even nominal speed-up is 1.74 times for $\varphi = 0$, 2.13 times for $\varphi = 0.2$ and 3.43 times for $\varphi = 0.4$ (Fig. 2). The definition of φ matters: at $\varphi = 0.4$, break-even is about 2.9 times if φ scales ρ, about 2.2 times if it scales $(\rho - 1)$, and about 3.4 times if it scales the fraction of experiments avoided, $(1 - 1/\rho)$. The last is the definition adopted in Eq. (2). Assuming $\varphi = 0$ therefore makes the required acceleration look much smaller than it is.

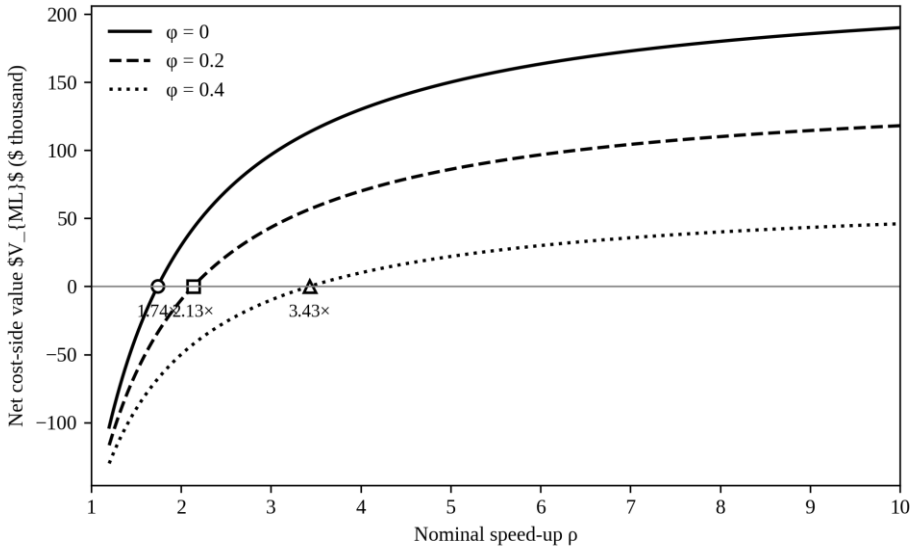


Figure 2. Net cost-side value of machine learning against nominal speed-up for three levels of model discrepancy φ. Markers show break-even. Hypothetical inputs: N_0 = 200, c_{exp} = \$2,000, Δt = 6 months, R_m = \$30,000 per month, C_{ML} = \$350,000, δ = 1.

Table 4. Break-even nominal speed-up ρ_{BE} (hypothetical inputs of Section 3.3).

φ	Baseline (δ = 1)	Discount and risk (δ = 0.5)	Δt also eroded, $\Delta t_{eff} = (1 - \varphi)\Delta t$
0	1.74×	2.86×	1.74×
0.1	1.89×	3.60×	2.09×
0.2	2.13×	5.33×	2.81×
0.3	2.55×	14.0×	5.00×
0.4	3.43×	No finite solution	No finite solution
0.5	6.67×	No finite solution	No finite solution
0.6	No finite solution	No finite solution	No finite solution

The existence limits on φ are 0.575 (baseline), 0.35 (δ = 0.5) and 0.397 (eroded Δt). Discounting and risk matter as much as discrepancy: at δ = 0.5, break-even at φ = 0.2 more than doubles.

Worked example. A venture tunes a lattice with five design parameters. A conventional campaign (a coarse grid) needs N_0 = 200 experiments at \$2,000 each, or \$400,000. Active learning with a Gaussian-process surrogate selects the next specimen where uncertainty is highest. In a published autonomous-experimentation study on additively manufactured polymer structures, this approach cut the experiments needed by about 60-fold relative to a grid search [28]; that is a different, macroscale system, so a far more modest ρ = 4 is assumed. The ideal campaign needs 50 experiments. If φ = 0.2, then 0.2 × 150 = 30 of the avoided experiments are performed after all, giving 80 experiments and ρ_{eff} = 200/80 = 2.5. The saving is 120 experiments, or \$240,000, before the cost C_{ML} and before any market-entry benefit.

3.4. Learning-curve and runway results

Table 5 reports the cost threshold n^*, time-to-volume, cumulative production cost, zero-margin volume and peak net outlay for three learning rates. The target is reachable at a learning rate of 20% (about 1,300 parts), marginal at 15% (about 18,000 parts and 18 months at the assumed throughput)

and out of reach at 10% (3.8 million parts, which is centuries of output; Fig. 3). The result is highly sensitive to b, which is itself poorly known for this technology.

Table 5. Learning-curve scenarios (c_1 = \$1,000; p = \$150; m_u = \$50; KQ = 1,000 parts per month; hypothetical).

L	b	n^* (parts)	T_{vol}	$C_{cum}(n^*)$	n_0 (parts)	L_{peak}
20%	0.80	1,277	1.3 months	\$188 thousand	362	\$26 thousand
15%	0.85	18,409	18.4 months	\$2.40 million	3,266	\$150 thousand
10%	0.90	3.79 million	about 316 years	\$447 million	263,234	\$7.1 million

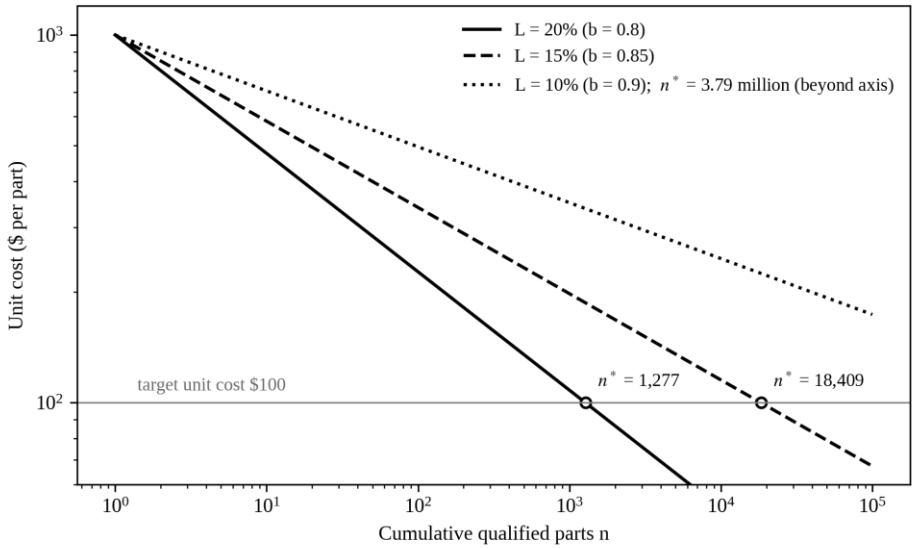


Figure 3. Unit cost against cumulative output for three learning rates (Wright's relationship). Markers show the cumulative volume n^* at which cost reaches the target of \$100. Hypothetical input: c_1 = \$1,000; the horizontal axis stops at 10^5 parts.

Applying the runway test with $\rho = 4$, $\varphi = 0.2$ (so $\rho_{eff} = 2.5$), C_{ML} = \$350,000, F = \$500,000, T_{dev} = 6 months, R_w = 24 months and K_{avail} = \$3.0 million gives the following. At L = 20%, the time is 6 + 1.3 = 7.3 months and K_{req} = \$1.2 million, so both conditions pass. At L = 15%, the time is 24.4 months (fails by 0.4 month) and K_{req} = \$3.41 million (fails). At L = 10%, both fail by orders of magnitude. In this example the learning rate decides the outcome far more than the value of machine learning in Eq. (3).

3.5. Uncertainty propagation results

With $\delta = 1$, $V_{ML} > 0$ in 64% of draws (median +\$61 thousand; 10th to 90th percentile −\$138 thousand to +\$297 thousand). With δ sampled, $V_{ML} > 0$ in 50% of draws (median about −\$1 thousand). The volume n^* can be reached within 24 months in only 23% of draws, and $V_{ML} > 0$ and $T_{vol} \leq 24$ months hold together in 12%. Under these assumptions the production-volume condition, not the experimental-savings condition, is the binding constraint. This result is conditional on the assumed distributions and should be recomputed from a venture's own pilot data. It demonstrates the mechanics of uncertainty propagation and does not establish a real-world success probability.

3.6. Gate criteria for the readiness framework

Table 6 gives indicative pass/fail thresholds and the approximate technology readiness level for each gate. The thresholds are starting points to be calibrated to the application and customer; they are not standards. The sample-size guidance for D1 follows a rule of thumb from computer experiments [46] applied here as a starting point.

Table 6. Readiness gates, indicative pass thresholds and approximate technology readiness level (TRL).

Gate	Question	Indicative pass threshold	TRL
D1 Data	Is there a curated, versioned dataset linking geometry, process and measured properties?	At least 95% of records with complete geometry, process and test metadata; replicate scatter (coefficient of variation) of the key property at most 15%; at least about 10 physical specimens per design-space dimension for the initial surrogate	2–3
D2 Model validity	Do predictions, with uncertainty, hold on unseen physical tests?	On a locked held-out set of at least 20 specimens covering at least 5 distinct designs: median absolute percentage error at most 15% for the primary property; 90% prediction intervals covering 85–95% of measurements; upper 90% confidence bound on φ below the decision threshold φ_{max} (see text)	3–4
D3 Producibility	Can the designs be made to tolerance with repeatability?	Process capability index Cpk at least 1.33 on critical dimensions; batch-to-batch coefficient of variation of the target property at most 10% over at least 3 batches; at least 70% of parts passing inspection; knockdown factor stable within ±10% across batches	5–6
D4 Economics	Does cost per part reach the target price along a credible learning curve, within the runway?	Conditions of Eqs. (10) and (11) hold at the cautious learning rate, with T_{vol} at most $0.5R_w$ and K_{req} at most $0.8K_{avail}$; modeled unit cost at pilot volume at most 1.5 times the target	6–7
D5 Qualification	Is performance independently verified?	Test report from an accredited laboratory (for example ISO/IEC 17025) on at least 3 batches; measured properties within ±10% of claimed values; at least one customer acceptance test passed	7–8
D6 IP and data protection	Is the advantage protectable?	Priority filing made, or a documented trade-secret regime in place, for the architecture, process and dataset; freedom-to-operate review finds no blocking claims; data rights and licence terms clear	Parallel track from about TRL 3

The D2 threshold for φ should be set from the break-even condition and not as a universal constant: a venture may take $\varphi_{max} = 1 - \theta$ for the baseline model, with a stricter bound when discounting, risk or delay erosion is included. This avoids treating one arbitrary cut-off as suitable for every product. A venture can use the gates to sequence spending: D1–D2 before scaling compute, D3 before committing to capacity, and D4–D6 before pursuing large customers.

3.7. Strategic implications

For founders. Treat experimentally validated data and calibrated models as important assets; competitors may find them harder to reproduce than a single geometry [34]. The results of Section 3.4 suggest that learning rate and throughput can matter more than the speed-up of the design search, so manufacturing-aware design deserves as much attention as machine-learning-assisted search. Initial markets should favor applications that can tolerate relatively high early unit costs, such as specialist components, protective equipment and research instrumentation, while production processes improve.

For investors. Request evidence at each gate, particularly validation on experiments not used for training and a defensible estimate of φ. Simulated record performance should be treated cautiously until model validity and manufacturability have been demonstrated at gates D2 and D3. Capital tranches can be linked to these milestones, and the runway test can be applied to each tranche.

For researchers and standards bodies. Shared benchmark datasets with experimental ground truth, consistent reporting of fabrication parameters and common validation protocols could reduce duplicated effort and help distinguish genuine machine-learning-enabled progress from gains caused by other factors. Reporting the fraction of screened candidates that failed physical validation would allow φ to be estimated for the first time.

3.8. Limitations and future research

The framework and the economic model are conceptual and have not been tested against venture case studies or real pilot data. This affects interpretation and generalizability in four ways. First, every numerical result (break-even speed-ups of 1.7–3.4 times, the 12–64% probabilities of Section 3.5 and the runway outcomes of Table 5) follows from hypothetical inputs, so the results show how the model behaves, not what a venture should expect; they should be read as conditional statements of the form "if φ and b take these values, then…" and not as forecasts. Second, the structural assumptions are untested: a constant cost per experiment, independence between ρ and φ, a single risk factor δ, a Wright-type learning curve in a low-throughput setting, constant throughput and a constant monthly margin. A different structure could change the conclusions, in particular the finding that production volume is the binding constraint. Third, the model is built for venture-stage, low-volume, specialist components with an exogenous price. It may not transfer to ventures whose value comes from a performance premium, from licensing or platform models, from high-volume production, or from established firms for which laboratory capacity is a sunk cost. It also values only cost-side effects, so it should not be used to argue that machine-learning-driven cost savings outweigh the value of a performance gain. Fourth, the literature is a limited basis: published work on machine-learning-designed nanostructures is dominated by simulation-led studies, fewer designs are fabricated and tested, and positive results are more likely to be published. No estimate of φ was found, and published successes probably understate it, so scenarios with low φ should be treated as optimistic. The review is also not a systematic search, so relevant studies may have been missed.

Future research should (a) create benchmark datasets pairing designs with measured properties; (b) estimate φ across fabrication routes and report it routinely; (c) estimate learning rates and throughput for nano- and microscale fabrication from pilot-line data; (d) calibrate and update φ, b and δ with Bayesian methods as a venture accumulates data, and allow dependence between parameters; (e) extend the model with market-side variables such as performance premium and willingness to pay; (f) combine multi-fidelity models with cost-aware autonomous experimentation; (g) include life-cycle and sustainability costs; and (h) test the framework on other classes of nanomaterials and on venture case studies.

4. Conclusions

Machine learning can shorten the nanomechanical design cycle, but commercial viability depends on three conditions: models must be calibrated and explicit about uncertainty, designs must respect manufacturing constraints, and production must have a credible route to a competitive unit cost

within the venture's runway. The Design-to-Venture Readiness framework and the techno-economic model make these conditions easier to state and test. In the illustrative analysis, the discrepancy fraction φ raises the required speed-up substantially (from 1.7 to 3.4 times at $\varphi = 0.4$), and the learning rate and throughput of production matter more than the speed of the design search. The practical question is therefore not only whether machine learning can identify a better architecture, but whether it can reduce the cost and time needed to produce and qualify one. The model values only cost-side effects and uses hypothetical parameters; ventures should support any claim with experimental, economic and market evidence of their own.

References

[1] Meza, L. R., Das, S., & Greer, J. R. (2014). Strong, lightweight, and recoverable three-dimensional ceramic nanolattices. *Science*, *345*(6202), 1322–1326. https://doi.org/10.1126/science.1255908

[2] Bauer, J., Schroer, A., Schwaiger, R., & Kraft, O. (2016). Approaching theoretical strength in glassy carbon nanolattices. *Nature Materials*, *15*, 438–443. https://doi.org/10.1038/nmat4561

[3] Crook, C., Bauer, J., Guell Izard, A., Santos de Oliveira, C., Martins de Souza e Silva, J., Berger, J. B., & Valdevit, L. (2020). Plate-nanolattices at the theoretical limit of stiffness and strength. *Nature Communications*, *11*, 1579. https://doi.org/10.1038/s41467-020-15434-2

[4] Portela, C. M., Edwards, B. W., Veysset, D., Sun, Y., Nelson, K. A., Kochmann, D. M., & Greer, J. R. (2021). Supersonic impact resilience of nanoarchitected carbon. *Nature Materials*, *20*, 1491–1497. https://doi.org/10.1038/s41563-021-01033-z

[5] Bauer, J., Meza, L. R., Schaedler, T. A., Schwaiger, R., Zheng, X., & Valdevit, L. (2017). Nanolattices: An emerging class of mechanical metamaterials. *Advanced Materials*, *29*, 1701850. https://doi.org/10.1002/adma.201701850

[6] Jiao, P., Mueller, J., Raney, J. R., Zheng, X., & Alavi, A. H. (2023). Mechanical metamaterials and beyond. *Nature Communications*, *14*, 6004. https://doi.org/10.1038/s41467-023-41679-8

[7] Deshpande, V. S., Ashby, M. F., & Fleck, N. A. (2001). Foam topology: Bending versus stretching dominated architectures. *Acta Materialia*, *49*(6), 1035–1040. https://doi.org/10.1016/S1359-6454(00)00379-7

[8] Gibson, L. J., & Ashby, M. F. (1997). *Cellular solids: Structure and properties* (2nd ed.). Cambridge University Press.

[9] Greer, J. R., & De Hosson, J. T. M. (2011). Plasticity in small-sized metallic systems: Intrinsic versus extrinsic size effect. *Progress in Materials Science*, *56*(6), 654–724. https://doi.org/10.1016/j.pmatsci.2011.01.005

[10] Schaedler, T. A., Jacobsen, A. J., Torrents, A., Sorensen, A. E., Lian, J., Greer, J. R., Valdevit, L., & Carter, W. B. (2011). Ultralight metallic microlattices. *Science*, *334*(6058), 962–965. https://doi.org/10.1126/science.1211649

[11] Zheng, X., Lee, H., Weisgraber, T. H., Shusteff, M., DeOtte, J., Duoss, E. B., Kuntz, J. D., Biener, M. M., Ge, Q., Jackson, J. A., Kucheyev, S. O., Fang, N. X., & Spadaccini, C. M. (2014). Ultralight, ultrastiff mechanical metamaterials. *Science*, *344*(6190), 1373–1377. https://doi.org/10.1126/science.1252291

[12] Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., & Walsh, A. (2018). Machine learning for molecular and materials science. *Nature*, *559*, 547–555. https://doi.org/10.1038/s41586-018-0337-2

[13] Bastek, J.-H., & Kochmann, D. M. (2023). Inverse design of nonlinear mechanical metamaterials via video denoising diffusion models. *Nature Machine Intelligence*, *5*, 1466–1475. https://doi.org/10.1038/s42256-023-00762-x

[14] Song, J., Lee, J., Kim, N., & Min, K. (2024). Artificial intelligence in the design of innovative metamaterials: A comprehensive review. *International Journal of Precision Engineering and Manufacturing*, *25*(1), 225–244. https://doi.org/10.1007/s12541-023-00857-w

[15] Turysbekov, G., Auyeskhan, U., Yankin, A., Perveen, A., & Talamona, D. (2026). A review of machine learning applications in mechanical metamaterial design. *Materials*, *19*(13), 2766. https://doi.org/10.3390/ma19132766

[16] He, L., Li, Y., Torrent, D., Zhuang, X., Rabczuk, T., & Jin, Y. (2023). Machine learning assisted intelligent design of meta structures: A review. *Microstructures*, *3*(4), 2023037. https://doi.org/10.20517/microstructures.2023.29

[17] Zhang, C., & Zhao, Y. F. (2024). A critical review on the application of machine learning in supporting auxetic metamaterial design. *Journal of Physics: Materials*, *7*(2), 022004. https://doi.org/10.1088/2515-7639/ad33a4

[18] Gu, G. X., Chen, C.-T., Richmond, D. J., & Buehler, M. J. (2018). Bioinspired hierarchical composite design using machine learning: Simulation, additive manufacturing, and experiment. *Materials Horizons*, *5*, 939–945. https://doi.org/10.1039/C8MH00653A

[19] Bessa, M. A., Bostanabad, R., Liu, Z., Hu, A., Apley, D. W., Brinson, C., Chen, W., & Liu, W. K. (2017). A framework for data-driven analysis of materials under uncertainty: Countering the curse of dimensionality. *Computer Methods in Applied Mechanics and Engineering*, *320*, 633–667. https://doi.org/10.1016/j.cma.2017.03.037

[20] Bastek, J.-H., Kumar, S., Telgen, B., Glaesener, R. N., & Kochmann, D. M. (2022). Inverting the structure–property map of truss metamaterials by deep learning. *Proceedings of the National Academy of Sciences*, *119*(1), e2111505119. https://doi.org/10.1073/pnas.2111505119

[21] Kumar, S., Tan, S., Zheng, L., & Kochmann, D. M. (2020). Inverse-designed spinodoid metamaterials. *npj Computational Materials*, *6*, 73. https://doi.org/10.1038/s41524-020-0341-6

[22] Bendsøe, M. P., & Sigmund, O. (2003). *Topology optimization: Theory, methods and applications*. Springer.

[23] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). 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. https://doi.org/10.1016/j.jcp.2018.10.045

[24] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, *3*, 422–440. https://doi.org/10.1038/s42254-021-00314-5

[25] Behler, J., & Parrinello, M. (2007). Generalized neural-network representation of high-dimensional potential-energy surfaces. *Physical Review Letters*, *98*, 146401. https://doi.org/10.1103/PhysRevLett.98.146401

[26] Deringer, V. L., Caro, M. A., & Csányi, G. (2019). Machine learning interatomic potentials as emerging tools for materials science. *Advanced Materials*, *31*, 1902765. https://doi.org/10.1002/adma.201902765

[27] Lookman, T., Balachandran, P. V., Xue, D., & Yuan, R. (2019). Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*, *5*, 21. https://doi.org/10.1038/s41524-019-0153-8

[28] Gongora, A. E., Xu, B., Perry, W., Okoye, C., Riley, P., Reyes, K. G., Morgan, E. F., & Brown, K. A. (2020). A Bayesian experimental autonomous researcher for mechanical design. *Science Advances*, *6*(15), eaaz1708. https://doi.org/10.1126/sciadv.aaz1708

[29] Vangelatos, Z., et al. (2021). Strength through defects: A novel Bayesian approach for the optimization of architected materials. *Science Advances*, *7*(41), eabk2218. https://doi.org/10.1126/sciadv.abk2218

[30] Auerswald, P. E., & Branscomb, L. M. (2003). Valleys of death and Darwinian seas: Financing the invention to innovation transition in the United States. *Journal of Technology Transfer*, *28*, 227–239. https://doi.org/10.1023/A:1024980525678

[31] Mankins, J. C. (1995). *Technology readiness levels: A white paper*. NASA Office of Space Access and Technology.

[32] Cooper, R. G. (1990). Stage-gate systems: A new tool for managing new products. *Business Horizons*, *33*(3), 44–54. https://doi.org/10.1016/0007-6813(90)90040-I

[33] Osterwalder, A., & Pigneur, Y. (2010). *Business model generation*. Wiley.

[34] Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, *15*(6), 285–305. https://doi.org/10.1016/0048-7333(86)90027-2

[35] Ries, E. (2011). *The lean startup*. Crown Business.

[36] Blank, S. (2013). Why the lean start-up changes everything. *Harvard Business Review*, *91*(5), 63–72.

[37] Page, M. J., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, *372*, n71. https://doi.org/10.1136/bmj.n71

[38] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems* (Vol. 30). arXiv:1612.01474

[39] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning* (PMLR 48, pp. 1050–1059). arXiv:1506.02142

[40] Kennedy, M. C., & O'Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B, 63*(3), 425–464. https://doi.org/10.1111/1467-9868.00294

[41] ISO/ASTM. (2021). *Additive manufacturing — General principles — Fundamentals and vocabulary* (ISO/ASTM 52900:2021). International Organization for Standardization / ASTM International.

[42] Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences, 3*(4), 122–128. https://doi.org/10.2514/8.155

[43] Nagy, B., Farmer, J. D., Bui, Q. M., & Trancik, J. E. (2013). Statistical basis for predicting technological progress. *PLoS ONE, 8*(2), e52669. https://doi.org/10.1371/journal.pone.0052669

[44] Argote, L., & Epple, D. (1990). Learning curves in manufacturing. *Science, 247*(4945), 920–924. https://doi.org/10.1126/science.247.4945.920

[45] Snapp, K. L., Verdier, B., Gongora, A. E., Silverman, S., Adesiji, A. D., Morgan, E. F., Lawton, T. J., Whiting, E., & Brown, K. A. (2024). Superlative mechanical energy absorbing efficiency discovered through self-driving lab-human partnership. *Nature Communications, 15*. https://doi.org/10.1038/s41467-024-48534-4

[46] Loeppky, J. L., Sacks, J., & Welch, W. J. (2009). Choosing the sample size of a computer experiment: A practical guide. *Technometrics, 51*(4), 366–376. https://doi.org/10.1198/TECH.2009.08040