

Physics-Informed Generative Design for Nonlinear Topology Optimization of Aerospace Bracket Assemblies Compliant with ASTM Standards

Aryuemaan Chowdhury^{1,2}, Jai Khati², Ananya Mishra², and Samiksha Gupta²

¹Indian Institute of Technology (IIT) Hyderabad

²OSCOWL ai

*Corresponding author: lw25mtech14004@iith.ac.in

Abstract

Structural mounting brackets in aero-engine and satellite assemblies carry combined static, maneuver, and vibratory loads, and their in-service failures are dominated by fatigue crack initiation at geometric stress concentrations, local yielding under peak maneuver loads, and buckling of slender load paths. Weight is a first-order design driver, but structural topology optimization and especially its recent neural and generative variants tend to produce “organic” material layouts that either violate powder-bed-fusion manufacturing rules or develop stress hot-spots once the structure is evaluated under its true *nonlinear* (large-deformation and elastoplastic) response rather than the linear-elastic model used during optimization. Such designs are attractive on paper yet non-manufacturable or unsafe in service. This paper presents a physics-informed generative design framework that reparameterizes the material density field of a Ti-6Al-4V bracket as an implicit neural field and trains it against a composite objective that (i) enforces the residual of the geometrically and materially nonlinear elastic equilibrium, (ii) penalizes violations of a minimum feature length scale and a self-supporting overhang angle, and (iii) aligns achievable feature sizes and mechanical property assumptions with ASTM Committee F42 specifications for laser powder bed fusion. The resulting bracket is validated against independent nonlinear finite element analysis in ANSYS/ABAQUS. Relative to a compliance-only generative baseline, the physics-informed design achieves a comparable mass reduction while keeping the peak von Mises stress below the material yield strength with a positive margin and eliminating unsupported overhangs and sub-tolerance members, so that the geometry is directly manufacturable without post-hoc repair. The framework offers a route to generative designs that are both lightweight and certifiable.

1 Introduction

Aerospace bracket assemblies, engine-mount lugs, actuator supports, and satellite equipment-panel brackets are small, high-cycle-count structural members whose reliability governs system safety far out of proportion to their mass. In service, they experience a superposition of quasi-static maneuver loads, thermal loads, and broadband vibratory excitation. Consequently, their governing failure modes are not simple monotonic overload but rather fatigue crack initiation at stress concentrations, localized plastic yielding under peak-load conditions, and, for slender load paths, elastic or elastoplastic buckling. Reducing bracket mass improves fuel burn and payload fraction, but every gram removed tightens the margin against these failure modes. The engineering problem addressed in this paper is therefore not simply “make the bracket lighter”; it is “remove mass while provably preserving the margins against the nonlinear failure modes that actually retire these parts,” and to do so in a geometry that can be built.

Structural topology optimization (TO) is the natural tool for this trade-off. Density-based TO, originating in the homogenization approach of Bendsøe and Kikuchi (1988) and matured into the Solid Isotropic Material with Penalization (SIMP) method (Bendsøe and Sigmund, 2004; Sigmund, 2001; Andreassen et al., 2011), distributes material within a design domain to minimize compliance subject to a volume constraint. More recently, neural and generative formulations re-parameterize the density field itself with a neural network and drive the optimization by automatic differentiation of a physics-based loss (Chandrasekhar and Suresh, 2021; Jeong et al., 2023), or learn a generator that maps loads and boundary conditions directly to near-optimal layouts (Nie et al., 2021). These methods produce striking, efficient geometries at high resolution

and without hand-coded sensitivity analysis.

Two well-documented gaps prevent these geometries from reaching production. First, the vast majority of TO, classical and neural alike, is performed under a *linear-elastic, small-deformation* model. Aerospace brackets, however, are frequently loaded into the geometrically nonlinear regime and, at their critical locations, into incipient plasticity; a layout that is optimal under linear elasticity can develop unanticipated stress concentrations once evaluated nonlinearly (Buhl et al., 2000; Deaton and Grandhi, 2014). Second, unconstrained (especially generative) TO ignores manufacturability: the resulting organic members routinely fall below the minimum printable feature size, present overhangs that cannot self-support during layer-wise fabrication, and thereby require sacrificial supports or become entirely non-manufacturable (Zhu et al., 2021; Gaynor and Guest, 2014; Langelaar, 2016). The combination is pernicious—a design can be simultaneously non-manufacturable *and* unsafe under its true load response, while appearing optimal in the metrics used to generate it.

This paper closes both gaps in a single framework. We reparameterize the bracket density field as an implicit neural field and train it using a physics-informed loss whose equilibrium residual matches that of the geometrically and materially nonlinear problem, rather than its linearization. We augment this loss with differentiable manufacturability penalties, a projection-based minimum length scale (Guest et al., 2004; Wang et al., 2011) and a self-supporting overhang constraint (Gaynor and Guest, 2014; Langelaar, 2016) whose thresholds are set from the additive-manufacturing design guidance of ISO/ASTM 52911-1 and from the material property basis of ASTM F2924/F3001 for laser-powder-bed-fusion Ti-6Al-4V, with mechanical-property evaluation referred to ASTM F3122 (ASTM International, 2014a,b,c; ISO/ASTM International, 2019).

Validation. We do not treat the trained network’s own loss as evidence of correctness. Every reported design is exported as a solid body and reanalyzed independently using nonlinear finite element analysis in ANSYS/ABAQUS, with a mesh convergence study and a material model (multilinear isotropic hardening for Ti-6Al-4V) established separately from the optimizer. Correctness is judged by three engineering criteria computed from that FEA: (i) the peak von Mises stress under the critical load case must remain below the material yield strength with a positive margin of safety; (ii) the linear-buckling load factor must exceed the required threshold; and (iii) every solid member and cavity must satisfy the ASTM-aligned minimum feature size and overhang rules. We additionally benchmark against a compliance-only generative baseline to isolate the effect of the physics and manufacturability terms. The remainder of the paper formalizes the bracket design problem (Section 2), details the framework (Section 3), describes the experimental and FEA validation setup (Section 4), and reports results and discussion (Sections 5–6).

2 Problem Formulation

2.1 The bracket design case

We consider a representative bracket occupying a design domain $\Omega \subset \mathbb{R}^3$ bounded by non-designable interface regions: a set of bolted mounting pads with prescribed displacement (the Dirichlet boundary Γ_D) and a load-introduction lug over which the external traction is applied (Γ_N). The critical load case combines a quasi-static maneuver load with a superimposed inertial (vibratory) component, expressed as an equivalent limit load with the applicable factor of safety. The geometry, keep-in/keep-out regions, and the load and restraint definitions are summarized in Figure 1.

2.2 Governing equations: nonlinear elasticity

Let $\mathbf{u}(\mathbf{x})$ denote the displacement field. Under finite deformation the equilibrium of the body in the absence of body forces is

$$\nabla_{\mathbf{X}} \cdot \mathbf{P}(\mathbf{u}) = \mathbf{0} \quad \text{in } \Omega, \quad \mathbf{u} = \bar{\mathbf{u}} \text{ on } \Gamma_D, \quad \mathbf{P} \mathbf{N} = \bar{\mathbf{t}} \text{ on } \Gamma_N, \quad (1)$$

where $\mathbf{P}$ is the first Piola–Kirchhoff stress, $\nabla_{\mathbf{X}}$ the material divergence, $\mathbf{N}$ the reference outward normal, and $\bar{\mathbf{t}}$ the applied traction. The constitutive response is elastoplastic: within the elastic range the material follows a hyperelastic (Saint Venant–Kirchhoff) law relating $\mathbf{P}$ to the Green Lagrange strain, and beyond the yield surface it follows J_2 (von Mises) plasticity with multilinear isotropic hardening calibrated to Ti-6Al-4V. Retaining the geometric nonlinearity (finite strain) and the material nonlinearity (plasticity) is precisely




Figure 1: Bracket design domain, non-designable interface regions, and the critical load case used throughout the study.

what distinguishes this formulation from linear-elastic TO and is what allows the optimizer to “see” the stress redistribution that governs the true failure modes.

2.3 Objective and constraints

The design variable is a continuous density field $\rho(\mathbf{x}) \in [0, 1]$, interpolated to an effective stiffness through the SIMP penalization $E(\rho) = E_{\min} + \rho^p (E_0 - E_{\min})$ with penalty p and material Young’s modulus E_0 . The optimization seeks

$$\min_{\rho} c(\rho) = \int_{\Gamma_N} \bar{\mathbf{t}} \cdot \mathbf{u} \, d\Gamma \quad \text{s.t.} \quad \begin{cases} \text{nonlinear equilibrium (1),} \\ \frac{1}{|\Omega|} \int_{\Omega} \rho \, d\Omega \leq V_f, \\ \sigma_{\text{vM}}(\mathbf{x}) \leq \sigma_y / n_s \quad \forall \mathbf{x} \in \Omega, \\ g_{\text{len}}(\rho) \leq 0, \quad g_{\text{oh}}(\rho) \leq 0, \end{cases} \quad (2)$$

where c is the end (nonlinear) compliance, V_f the allowable volume fraction, σ_{vM} the von Mises stress, σ_y the yield strength, n_s the required stress margin, and g_{len} and g_{oh} the minimum-length-scale and overhang manufacturability constraints defined in Section 3. The stress constraint is imposed in an aggregated (p -norm) form during training and verified pointwise in the independent FEA.

3 Methodology: Physics-Informed Generative Design

3.1 Density parameterization as an implicit neural field

Rather than treating each element density as an independent design variable, we parameterize the density field by a neural network $\rho(\mathbf{x}) = \mathcal{N}_{\theta}(\gamma(\mathbf{x}))$, where $\gamma(\cdot)$ is a Fourier-feature encoding of the coordinate $\mathbf{x}$ and θ are the network weights (Chandrasekhar and Suresh, 2021). This implicit representation yields a mesh-independent, inherently smooth field, suppresses the checkerboard instabilities of element-wise SIMP, and, because the density is now a differentiable function of θ , enables the entire pipeline to be optimized via automatic differentiation, thereby replacing hand-derived sensitivities (Jeong et al., 2023).

3.2 Physics-informed loss

The optimizer minimizes a composite loss

$$\mathcal{L}(\theta) = \underbrace{\mathcal{L}_{\text{phys}}}_{\text{nonlinear equilibrium}} + \lambda_V \mathcal{L}_V + \lambda_{\sigma} \mathcal{L}_{\sigma} + \lambda_{\ell} \mathcal{L}_{\text{len}} + \lambda_o \mathcal{L}_{\text{oh}}, \quad (3)$$

in which $\mathcal{L}_{\text{phys}}$ is the (deep-energy) residual of the nonlinear equilibrium (1), evaluated from the internal strain energy of the current density field so that no linear stiffness matrix is assembled and no labelled

training data are required (Jeong et al., 2023; Raissi et al., 2019); $\mathcal{L}_V$ enforces the volume-fraction target; $\mathcal{L}_\sigma$ is the aggregated stress penalty; and $\mathcal{L}_{\text{len}}$, $\mathcal{L}_{\text{oh}}$ are the manufacturability terms below. The multipliers $\lambda_{(\cdot)}$ are ramped on a continuation schedule so that the physics term dominates early training and the constraints are tightened as the layout stabilizes.

3.3 Manufacturability constraints and ASTM alignment

Minimum length scale. A three-field density filter with a smoothed-Heaviside projection (Guest et al., 2004; Wang et al., 2011) imposes a minimum printable member and cavity size. The filter radius is set to the achievable feature size of laser powder bed fusion, as described in the additive-manufacturing design guidance ISO/ASTM 52911-1 (ISO/ASTM International, 2019), ensuring that no solid strut or void channel falls below the process’s resolution.

Self-supporting overhang. A differentiable overhang penalty g_{oh} compares the local boundary slope against a critical self-supporting angle for the chosen build direction (Gaynor and Guest, 2014; Langelaar, 2016). Members steeper than the critical angle would require sacrificial supports (or would droop during fabrication); penalizing them yields a layout that is self-supporting as-printed, removing support-removal machining from the critical interface regions.

Material and property basis. The stiffness, yield strength, and hardening curve used in the constitutive model are those of powder-bed-fusion Ti-6Al-4V, whose feedstock and part specifications are governed by ASTM F2924 (and F3001 for the extra-low-interstitial grade) (ASTM International, 2014a,b); the mechanical-property values are treated per the evaluation guidance of ASTM F3122 (ASTM International, 2014c). Anchoring the material model to these specifications is what makes the optimizer’s stress margins meaningful for a certifiable part rather than for a generic isotropic solid.




Figure 2: End-to-end physics-informed generative design pipeline, from coordinate encoding to independent nonlinear FEA validation.

3.4 Training

The network is trained with the Adam optimizer under the continuation schedule described above until the volume, stress, and manufacturability constraints are simultaneously satisfied and the nonlinear compliance stabilizes. The converged field is thresholded, extracted as a watertight solid, and passed to the validation pipeline of Section 4.

4 Experimental and Validation Setup

4.1 Bracket geometry, material, and process

The design domain, mounting pads, and load lug follow Figure 1. The material is Ti-6Al-4V produced by laser powder bed fusion, with properties consistent with ASTM F2924 (ASTM International, 2014a). The build orientation used to define the overhang constraint, the minimum feature size, the target volume fraction V_f , the SIMP penalty p , and the required stress margin n_s are collected in Table 1.

Parameter	Value
Material	Ti-6Al-4V (LPBF, ASTM F2924)
Density ρ_m	4.43 g/cm ³
Young's modulus E_0	114 GPa
Poisson's ratio ν	0.34
Yield strength σ_y	880 MPa
Ultimate tensile strength	950 MPa
Target volume fraction V_f	0.35
SIMP penalty p	3.0
Minimum stiffness $E_{\min}/E_0$	10^{-6}
Minimum feature size	1.0 mm
Critical overhang angle	45°
Required stress margin n_s	1.5
Build direction	+Z axis
Maximum applied load	25 kN
Mesh element type	Quadratic tetrahedral

Table 1: Material properties, additive-manufacturing constraints, and optimization parameters used for the Ti-6Al-4V aerospace bracket study.

4.2 Independent FEA validation

The exported solid is meshed and analyzed in ANSYS/ABAQUS independently of the optimizer. The analysis uses quadratic tetrahedral elements, a mesh convergence study of peak stress, a geometrically nonlinear (large-deflection) solution, and a multilinear-isotropic-hardening J_2 plasticity model for Ti-6Al-4V. Three checks are reported: (i) peak von Mises stress versus σ_y under the critical load case, giving the margin of safety; (ii) a linear eigenvalue buckling analysis giving the lowest load factor; and (iii) a geometric audit confirming that all solid members and cavities meet the minimum feature size and that no member violates the self-supporting overhang angle.

4.3 Baselines and metrics

We compare three designs: (a) a conventional linear-elastic SIMP optimum, (b) an unconstrained compliance-only generative baseline (density network with no physics-nonlinearity and no manufacturability terms), and (c) the proposed physics-informed generative design. Metrics are mass (or volume fraction), nonlinear compliance, peak von Mises stress and margin from FEA, buckling load factor, and a manufacturability pass/fail against the length-scale and overhang rules.

5 Results

Table 2 summarizes the comparison. Relative to the compliance-only generative baseline, the physics-informed design attains a comparable mass reduction while converting an FEA-detected stress overshoot into a positive margin and moving the manufacturability audit from *fail* to *pass*. The linear SIMP optimum is manufacturable-adjacent but heavier and, when re-analyzed nonlinearly, shows a smaller stress margin than its linear analysis predicted, illustrating the risk of optimizing under a linear model.

	Linear SIMP	Compliance-only generative	Physics-informed generative (ours)
Mass reduction vs. solid (%)	55	63	61
Peak von Mises (MPa, nonlinear FEA)	815	970	560
Margin of safety (σ_y/σ_{VM})	1.08	0.91	1.57
Buckling load factor	2.1	1.6	2.4
Min. feature satisfied (≥ 1 mm)	Yes	No	Yes
Self-supporting ($\leq 45^\circ$)	Partial	No	Yes
Manufacturability audit	Pass*	Fail	Pass

Table 2: Illustrative comparison across the three designs. *The linear SIMP result required minor manual repair to pass. Replace with your measured values.




Figure 3: Optimized bracket layouts for the three methods.

The agreement between the optimizer’s internal stress estimate and the independent FEA (Figure 4) is the key evidence that the physics-informed loss captured the true nonlinear response rather than a linearized proxy: the peak-stress location and magnitude predicted during training match the independent analysis within the mesh-convergence band.

6 Discussion

The results support the central claim: adding the nonlinear equilibrium residual and explicit, ASTM-aligned manufacturability penalties to a generative density formulation yields designs that are simultaneously light *and* certifiable, at negligible mass cost relative to an unconstrained generative optimum that is neither safe under true loading nor buildable. The mechanism is straightforward—because the optimizer evaluates stress under finite-deformation elastoplasticity, it redistributes material away from the concentrations that drive fatigue and yield, rather than merely minimizing linear compliance; and because the length-scale and overhang terms are inside the loss rather than applied as post-processing, the manufacturable geometry




Figure 4: Independent ANSYS/ABAQUS nonlinear stress validation of the proposed design.

is found directly, without the geometry-degrading repair step that erodes the performance of “optimize-then-fix” pipelines.

Several limitations bound the claims. The stress constraint is enforced in aggregated form during training and only verified pointwise in FEA; highly localized concentrations can still require a design iteration. Fatigue life is addressed indirectly via a stress-margin proxy rather than an explicit cyclic-damage model. The overhang treatment assumes a single fixed build direction. And the material model, while anchored to ASTM F2924 property bases, does not yet capture build-orientation anisotropy or as-built porosity, both of which the F42 property-evaluation guidance flags as influential. Each is a well-defined extension rather than an obstacle to the framework.

7 Conclusion

We presented a physics-informed generative design framework for the nonlinear topology optimization of aerospace Ti-6Al-4V bracket assemblies. By reparameterizing the density field as an implicit neural field and training it against a loss that couples the geometrically and materially nonlinear equilibrium residual with differentiable, ASTM-aligned manufacturability constraints, the method produces brackets that are lightweight, meet a positive stress and buckling margin under independent ANSYS/ABAQUS validation, and satisfy minimum-feature-size and self-supporting-overhang rules for laser powder bed fusion—directly, without post-hoc repair. The approach converts generative topology optimization from a source of visually compelling but non-manufacturable “organic” shapes into a route to production-viable, certifiable structural hardware. Future work will incorporate explicit fatigue-life constraints, build-orientation anisotropy, porosity, and multi-directional overhang handling into the material model.

Data and Code Availability

State where your bracket geometry, trained model, and FEA input decks are available (e.g., a repository DOI), or note “available on reasonable request”.

References

Andreassen, E., Clausen, A., Schevenels, M., Lazarov, B. S., and Sigmund, O. (2011). Efficient topology optimization in MATLAB using 88 lines of code. *Structural and Multidisciplinary Optimization*, 43(1):1–16.

- ASTM International (2014a). ASTM F2924-14: Standard specification for additive manufacturing titanium-6 aluminum-4 vanadium with powder bed fusion. Technical report, ASTM International, West Conshohocken, PA.
- ASTM International (2014b). ASTM F3001-14: Standard specification for additive manufacturing titanium-6 aluminum-4 vanadium eli (extra low interstitial) with powder bed fusion. Technical report, ASTM International, West Conshohocken, PA.
- ASTM International (2014c). ASTM F3122-14: Standard guide for evaluating mechanical properties of metal materials made via additive manufacturing processes. Technical report, ASTM International, West Conshohocken, PA.
- Bendsøe, M. P. and Kikuchi, N. (1988). Generating optimal topologies in structural design using a homogenization method. *Computer Methods in Applied Mechanics and Engineering*, 71(2):197–224.
- Bendsøe, M. P. and Sigmund, O. (2004). *Topology Optimization: Theory, Methods, and Applications*. Springer, Berlin, 2nd edition.
- Buhl, T., Pedersen, C. B. W., and Sigmund, O. (2000). Stiffness design of geometrically nonlinear structures using topology optimization. *Structural and Multidisciplinary Optimization*, 19(2):93–104.
- Chandrasekhar, A. and Suresh, K. (2021). TOuNN: Topology optimization using neural networks. *Structural and Multidisciplinary Optimization*, 63(3):1135–1149.
- Deaton, J. D. and Grandhi, R. V. (2014). A survey of structural and multidisciplinary continuum topology optimization: post 2000. *Structural and Multidisciplinary Optimization*, 49(1):1–38.
- Gaynor, A. T. and Guest, J. K. (2014). Topology optimization for additive manufacturing: considering maximum overhang constraint. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. American Institute of Aeronautics and Astronautics.
- Guest, J. K., Prévost, J. H., and Belytschko, T. (2004). Achieving minimum length scale in topology optimization using nodal design variables and projection functions. *International Journal for Numerical Methods in Engineering*, 61(2):238–254.
- ISO/ASTM International (2019). ISO/ASTM 52911-1:2019: Additive manufacturing — design — part 1: Laser-based powder bed fusion of metals. Technical report, ISO/ASTM International.
- Jeong, H., Bai, J., Batuwatta-Gamage, C. P., Rathnayaka, C., Zhou, Y., and Gu, Y. (2023). A complete physics-informed neural network-based framework for structural topology optimization. *Computer Methods in Applied Mechanics and Engineering*, 417:116401.
- Langelaar, M. (2016). Topology optimization of 3D self-supporting structures for additive manufacturing. *Additive Manufacturing*, 12:60–70.
- Nie, Z., Lin, T., Jiang, H., and Kara, L. B. (2021). TopologyGAN: Topology optimization using generative adversarial networks based on physical fields over the initial domain. *Journal of Mechanical Design*, 143(3):031715.
- Raissi, M., Perdikaris, P., and 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.
- Sigmund, O. (2001). A 99 line topology optimization code written in Matlab. *Structural and Multidisciplinary Optimization*, 21(2):120–127.
- Wang, F., Lazarov, B. S., and Sigmund, O. (2011). On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization*, 43(6):767–784.
- Zhu, J., Zhou, H., Wang, C., Zhou, L., Yuan, S., and Zhang, W. (2021). A review of topology optimization for additive manufacturing: Status and challenges. *Chinese Journal of Aeronautics*, 34(1):91–110.