

CONTEXT-AWARE DEEP REINFORCEMENT LEARNING FOR MULTI-OBJECTIVE TOPOLOGY OPTIMIZATION: FROM GLOBAL COMPLIANCE TO LOCAL STRESS CONSTRAINTS

Valentin Würz, Prof. Dr. Christian Weißenfels

Augsburg University, Institute of Materials Resource Management

Keywords: Structural Topology Optimization, Deep Reinforcement Learning, JAX-FEM, Stress-based Design, Load Generalization.

Abstract

Conventional topology optimization (TO) frameworks, such as the Solid Isotropic Material with Penalization (SIMP) method, have been highly successful in minimizing global structural compliance. However, these methods often rely on the differentiability of the objective function and struggle with local, non-smooth constraints, most notably local maximum stress limits, due to the singularity phenomenon and high sensitivity fluctuations. We propose a hybrid structural optimization strategy that replaces traditional gradient-driven update rules with a Deep Reinforcement Learning (DRL) controller, enabling the direct optimization of structures under non-differentiable and local mechanistic constraints.

Our framework parameterizes designs through a compact graph of geometric degrees of freedom. A set of nodes or articulation points defines the connectivity and support for an implicit truss-like structure, which is then mapped to a high-resolution voxelized density field. Structural performance is evaluated using a high-performance finite-element (FE) solver (JAX-FEM). Unlike traditional gradient-based density methods, our Soft Actor-Critic (SAC) agent iteratively modifies the node positions and local radii based on a multi-objective reward engine. This approach allows for the seamless integration of local maximum stress constraints, moving beyond global compliance to prevent localized material failure.

Crucially, this method addresses three major bottlenecks in modern design optimization:

Scalability to 3D: By leveraging the efficiency of JAX-FEM and a compact graph representation reducing degrees of freedom, the framework scales to complex 3D domains where traditional element-wise density updates become computationally expensive.

Generalization: Once trained, the RL policy demonstrates zero-shot generalization, producing high-performance topologies for unseen load cases and boundary conditions without requiring re-optimization.

Complex objectives: The reward-based formulation allows for the inclusion of non-differentiable manufacturing constraints and stress-based objectives that are typically infeasible due to inefficient scaling of elementwise constraint formulations. Results demonstrate that the learned policy captures fundamental structural principles, resulting in high-stiffness topologies that strictly respect volume budgets and local stress thresholds while maintaining high flexibility in the design of variable customizable objective functions. This work bridges the gap between geometric design and mechanistic performance, offering a robust alternative to traditional TO for multi-constraint engineering applications.