Abstract

Structural topology optimization seeks the optimal distribution of material within a design domain under equilibrium constraints and prescribed performance objectives. In density-based methods such as the Solid Isotropic Material with Penalization (SIMP) approach, the dominant computational cost arises from the repeated solution of large finite element systems at every design iteration. This burden becomes especially severe at engineering-relevant resolutions and in three dimensions, limiting the practical use of high-fidelity topology optimization in design exploration and repeated analysis.

This dissertation develops a runtime-verified multifidelity framework for minimum-compliance topology optimization under a global volume constraint in linear elasticity. The central contribution is a Multifidelity Acceptance Safeguard (MFAS) that periodically evaluates coarse-mesh candidate designs on the fine mesh during the optimization loop. Candidate updates are accepted only if they satisfy the safeguard; otherwise, they are rejected and the coarse state is resynchronized from the last accepted fine-mesh design. By embedding verification within the runtime loop rather than applying it only after optimization, the framework limits the accumulation of coarse-model error while preserving fine-mesh reliability.

Within this verified framework, machine learning is incorporated as an optional warm-start mechanism rather than a replacement for physics-based optimization. A convolutional encoder-decoder (U-Net) predicts coarse-mesh initializations from geometry and boundary-condition descriptors, and all learned proposals remain subject to the same acceptance-rejection protocol as coarse-only candidates.

Four method variants are studied: a fine-mesh reference method, an unverified coarse-only method, a verified multifidelity method, and a verified multifidelity method with U-Net warm start. The framework is evaluated on representative two- and three-dimensional benchmark problems, and structured ablation studies examine the influence of verification frequency and acceptance tolerance. Paired-seed experiments and nonparametric tests are used to assess observed differences in compliance and runtime. The results support the thesis that acceleration in topology optimization becomes substantially more trustworthy when it is coupled with explicit in-the-loop verification.