A! Aalto University
School of Engineering

A Comparison of Topology Optimization Methods in Ansys Mechanical

Race-Car Suspension Bracket Case Study (Topology Optimization for Lightweight Engineering Structures)

Ijtihed Kilani
Joakim Grönroos

Aalto University, School of Engineering

Keywords: topology optimization, Ansys Mechanical, SIMP, level set, compliance minimization, suspension bracket

Advisor: Rauf Talha

Table of Contents

A Comparison of Topology Optimization Methods:

Physical Race-Car Suspension Bracket Case Study

(*Topology Optimization for Lightweight Engineering Structures*)

Ijtihed Kilani[1], Joakim Grönroos[1]

Group 5

[1]Aalto University, School of Engineering

April 3, 2026

Abstract

Topology optimization can produce large stiffness or weight improvements, but results are often difficult to compare across studies because software defaults, filters, and stopping criteria differ, also because manufacturability constraints are frequently applied only after optimization. This paper presents a controlled comparison of topology-optimization methods in Ansys Mechanical with a fixed problem definition and a fixed physical length scale, so that observed differences reflect the method and solver behavior rather than mesh artifacts. To accomplish this, two studies were made. First, a minimal pseudo 2-D benchmark plate is defined to document the full numerical workflow and to report convergence, final mass or volume fraction, and strain energy (equivalently compliance under linear statics). Second, the same reporting protocol is applied to a real suspension bracket from a race-car assembly, with non-design bolt interfaces and three representative load cases. Density-based SIMP and the built-in level-set method are emphasized, and additional Ansys methods are included on the benchmark to broaden context. Across both cases, level-set tends to yield sharper, CAD-ready boundaries and slightly lower mass, while SIMP achieves the lowest strain energy and stress measures under matched constraints. The paper provides reproducible setup details, matched metrics, and qualitative manufacturability observations to support fair side-by-side evaluation.

Keywords: topology optimization, SIMP, level-set, Ansys Mechanical, compliance, strain energy, minimum member size, manufacturability, additive manufacturing, benchmark study

1 Introduction

Topology optimization aims to produce lighter, stiffer, or better-cooled parts that can be made in the real world. In practice, many simulated shapes are hard or costly to manufacture, and results can vary across tools because defaults, filters, and optimizers differ [1,5]. Manufacturability is treated as a first-class goal and fair, side-by-side comparisons on identical cases are used. Our focus is placed on SIMP in Ansys, and a level-set run is included to check how a sharp boundary affects geometry quality and printability [1,2,3].

Background and historical context. Topology optimization grew from structural mechanics and variational design. The density (SIMP) approach became the practical workhorse and was popularized by Bendsoe and Sigmund, together with simple reference codes that clarified filters and continuation [8,9]. Reliable large-scale optimization was enabled by MMA (Method of Moving Asymptotes) [11]. Level-set methods brought sharp interfaces and PDE-based boundary motion into the field, with significantly clean handoff to CAD/Printing [12]. AM (Additive manufacturing) then pushed methods to include overhang and minimum thickness during the run rather than after it, with projection and AM-aware filters becoming slowly more common [6,7,10].

Manufacturability in scope. Real parts must be made. For AM, build direction, overhang, and minimum feature size would be considered, while for machining or casting, draw direction, tool access, and practical radii are considered [5,6,7]. Prior studies often mix software releases and unclear settings, which makes results quite hard to compare [1,3,4,5]. Geometry, loads, supports, target volume, mesh family and size, and a physical length scale are fixed to isolate method and software effects.

Initial goals. A small, reproducible benchmark is built and practical limits are respected:

- A benchmark structural case with fixed geometry, loads, supports, volume fraction and mesh size is used for comparisons.
- SIMP and a level-set method are in focus while multiple methods are compared using the benchmark case.
- The effect of manufacturability (minimum size, overhang/draw constraints, tooling constraints) on performance metrics is briefly discussed.
- The key settings and parameters needed to reproduce results are disclosed and CAD-ready boundaries are shown.
- A real world design case is used as an example to simplify visualization of the problem, and to further validate the results obtained in the benchmark case.

2 Literature Review

2.1 Current publications that are used in this work

Six of the reviewed publications that match the tools and cases are used. The Ansys bracket paper runs SIMP with a volume cap and reports stiffness-per-mass gains on sheet-metal-like parts [1]. The IOP paper runs SIMP in Ansys with a parametric check and a response surface and confirms the steps on a newer release [3]. The robot limb paper applies SIMP in Ansys to a complex limb and reports weight cuts and stiffness checks, with some mismatch between tests and simulation [4]. The CMES paper builds a level-set method in a commercial multiphysics solver with a stable update and shows smooth boundaries and good compliance values [2]. The MM Science survey lists features and constraints across several commercial tools (e.g., Ansys, Abaqus with Tosca, Altair Inspire or OptiStruct, and nTop) and helps map which tool exposes which rule [5]. The manufacturing constraints paper explains filter- and projection-based length scale control and provides a model that can be reused [6].

2.2 Most important theories and results

This subsection first explains the main topology optimization formulations in general. After that, it describes how the relevant methods are actually implemented in Ansys Mechanical and how this differs from the more generic "optimization" tools in Ansys.

General formulations (literature view). In the literature, four main formulations are often discussed:

- density-based methods (SIMP),
- level-set methods,
- phase-field methods, and
- methods based on the topological derivative.

All of them solve roughly the same underlying problem: choose a layout of material inside a design region so that a structural objective (usually compliance) is minimized, subject to a volume constraint and possibly additional manufacturing rules.

Strain energy and compliance (objective used in this work). In linear elasticity, the work done by the external forces is stored in the structure as *strain energy*. Let u be the displacement field, $\varepsilon(u)$ the strain, and $\sigma(u)$ the stress. The (total) strain energy is

$$U(u) = \frac{1}{2}\int_{\Omega} \sigma(u) : \varepsilon(u)\, d\Omega.$$

For a linear elastic material, $\sigma = \mathbf{C} : \varepsilon$ with elasticity tensor $\mathbf{C}$, giving the equivalent form

$$U(u) = \frac{1}{2} \int_\Omega \varepsilon(u) : \mathbf{C} : \varepsilon(u) \, d\Omega.$$

In the finite-element setting, the same quantity is

$$U = \frac{1}{2} u^T K u,$$

where K is the global stiffness matrix.

For the common case of *load control* (fixed external load vector f) and small deformations, the equilibrium equation is $Ku = f$. The compliance is defined as

$$C = f^T u.$$

Using equilibrium, one obtains the standard identity

$$C = u^T K u = 2U.$$

Therefore, minimizing compliance is equivalent to minimizing strain energy (up to the constant factor 2) and produces the stiffest design for a fixed load and volume fraction. In this paper, the reported "strain energy" values from Ansys correspond to U, while the classical topology-optimisation objective is often written as compliance C; the two are directly comparable through $C = 2U$ in the linear static cases considered here.

SIMP (density-based). In SIMP, each finite element in the design domain has a design variable $\rho \in [0, 1]$ that represents how much material is present. The material stiffness is interpolated as

$$E(\rho) = E_{\min} + \rho^p \, (E_0 - E_{\min}),$$

where E_0 is the Young's modulus of the solid material, $E_{\min} > 0$ is a small "ersatz" stiffness for void elements, and p is a penalization factor. For a given density field, the state equation and compliance objective are

$$K(\rho) \, u = f, \qquad C(\rho) = f^T u,$$

with a volume constraint

$$\frac{1}{|\Omega|} \int_\Omega \rho \, d\Omega \leq \bar{v}.$$

Small values of p give gray layouts, while larger p drive the solution toward black–white designs [1,3,4,8,10].

Filters and projections. Without any smoothing, SIMP tends to produce checkerboard patterns and unrealistically thin members. A standard cure is to first average densities locally and then apply a smooth projection. A typical filter and projection looks like

$$\bar{\rho}_i = \frac{\sum_{j\in N_i} w_{ij}\,\rho_j}{\sum_{j\in N_i} w_{ij}}, \quad w_{ij} = \max\bigl(0,\, r_{\min} - \|x_i - x_j\|\bigr),$$

$$\tilde{\rho}_i = \frac{\tanh(\beta\eta) + \tanh\bigl(\beta(\bar{\rho}_i - \eta)\bigr)}{\tanh(\beta\eta) + \tanh\bigl(\beta(1-\eta)\bigr)}.$$

Here $r_{\min}$ sets a minimum member size, β controls how sharp the transition is, and $\eta \approx 0.5$ is a threshold. Increasing β during the run makes the final design more clearly black–white and less mesh-dependent [6,9].

Level-set methods. Level-set methods represent the structure by a level-set function $\phi(x)$:

$$\Omega_s = \{x \mid \phi(x) > 0\}, \quad \Omega_v = \{x \mid \phi(x) < 0\}, \quad \Gamma = \{x \mid \phi(x) = 0\}.$$

The interface moves according to an evolution law

$$\partial_t \phi = -V_n\,\|\nabla\phi\| + \tau\nabla^2\phi,$$

where V_n is the normal velocity derived from shape sensitivities, and τ is a smoothing parameter. Volume is enforced through a smoothed indicator,

$$\int_\Omega H_\beta(\phi)\,d\Omega \le \bar{v}\,|\Omega|.$$

This yields sharp, CAD-ready boundaries tha t are well suited for direct geometry export [2,12].

Phase-field and topological derivative. Phase-field methods use an order pa rameter $\phi \in [0,1]$ together with a perimeter penalty to control feature size, often by adding a Ginzburg–Landau term to the compliance [10]. Topological-derivative methods compute the change in objective when a tiny hole is inserted at a point and then add or remove material based on this derivative [10]. Both are powerful for research, but they are not the main focus in our study.

Ansys Mechanical implementation (methods used in this work). The version of Ansys Mechanical used in this work exposes two topology optimization methods:

- a density-based topology optimizer based on SIMP, and
- a level-set-based topology optimizer with a sharp solid–void boundary.

Phase-field and topological-derivative formula tions are not directly available and are therefore included only in the general literature review.

Density-based topology optimization in Ansys (step by step). In Ansys, the SIMP-based topology optimization roughly proceeds as follows:

1. The user defines a design region (the elements where material can be removed) and sets a target volume fraction (for example 25% of the original volume).

2. Ansys assigns a design variable ρ_e to each element in the design region and uses a SIMP-type law

$$E(\rho_e) = E_{\min} + \rho_e^p \, (E_0 - E_{\min})$$

   internally. The user does not usually se e the exact values of p or $E_{\min}$.

3. A finite element analysis is solved to obtain the displacement field u and the compliance $C = f^T u$ for the current density distribution.

4. Sensitivities of the objective and constraints with respect to the design variables are computed internally.

5. Ansys applies internal filtering and projection to enforce the chosen minimum member size. The user typically only specifies a target feature size; the tool chooses filter radius and projection steepness so that elements thinner than this size are di scouraged.

6. An internal gradient-based update (similar in spirit to OC/MMA) changes the densities, while keeping the volume fraction close to the target and respecting move limits for stability.

7. Steps 3–6 are repeated until the change in compliance is small or a maximum number of iterations is reached.

The user mostly interacts with high-level settings (design region, volume fraction, minimum member size, symmetry, and sometimes overhang rules), while Ansys handles the low-level numerical details.

Level-set topology optimization in Ansys (step by step). The Ansys level-set-based optimization follows the same physics but uses a moving boundary instead of densities:

1. The user defines a starting shape (for example a full design region with non-design regions cut out) and selects level-set topology optimization as the method.

2. Ansys represents the current shape by an internal level-set function $\phi(x)$, which marks solid and void regions on the mesh.

3. A finite element analysis is solved for the current shape, giving displacements and compliance for the active (solid) region.

4. Shape sensitivities are computed on the boundary. These are used to define a normal velocity V_n that tells the boundary where to move to reduce compliance while respecting the volume constraint.

5. The level-set field is updated according to an evolution law similar to

$$\partial_t \phi = -V_n \|\nabla \phi\| + \tau \nabla^2 \phi,$$

with internal steps for re-initialization and smoothing to keep ϕ well behaved.

6. Volume fraction and minimum member size are enforced by internal controls on how much the boundary is allowed to move and how fine features can become. The user again sets only higher-level values such as target volume and member size.

7. The analysis and update steps are repea ted until the objective and geometry stop changing significantly.

Compared to SIMoP, this produces sharper, more CAD-ready geometries but gives the user less freedom to tune the underlying numerical parameters.

Relation to standard optimization in Ansys. It is useful to contrast these topology tools with more classical optimization in Ansys. Size and shape optimization work on explicit parameters (for example cross-section areas, thicknesses, or spline control points) and use general-purpose optimizers. They cannot create new holes; they only adjust an existing geometry. In topology optimization, Ansys instead optimizes over fields (densities or a level-set function), can add or remove mateddrial, and has built-in filters for minimum feature size and manufacturability. This difference in design representation and solver is the main reason why the topology optimization results in this work look and behave differently from standard parameter optimizations in the same software.

2.3 Well-posedness and length-scale regularization

The basic compliance-minimization problem without any spatial regularization is ill-posed. Sequences of designs can keep forming ever finer features that reduce compliance, so the infimum is not attained and the result depends on the mesh. In practice this shows up as checkerboards, microstructures, and sensitivity to small meshing or boundary changes. The root cause is the lack of compactness of the admissible designs in the unregularized formulation [8,10].

There are two standard cures that introduce a fixed physical length scale. The first is *variational regularization*, where a perimeter or curvature term is added in the continuous problem. This fits naturally with level-set and phase-field formulations and promotes crisp, smooth boundaries with existence properties in the continuous setting [10,12]. The second is *operator regularization*, where design fields or their sensitivities are smoothed and then mapped toward 0–1 by a smooth projection. This is the common industrial approach with density methods and suppresses checkerboards while imposing a practical minimum feature size on the discrete grid [6,8,9,10].

With a fixed length scale the optimization becomes well posed in practice. Designs show mesh-independent trends, continuation produces near-binary layouts rather than gray fog, and loads and supports can be applied over small but finite pads without spurious singularities. This restores stability and makes comparisons meaningful across meshes and tools.

This study follows that principle. A single physical length scale is kept fixed across methods and software so that observed differences reflect formulation and solver rather than mesh artifacts. Level-set runs realize the length scale through curvature or perimeter control, while density-based runs realize it through filtering and smooth projection; both enforce the same idea of a physical, process-aware minimum feature size [2,6,10,12].

2.4 SIMP as relaxation and material interpolation

The original topology task is discrete: each point is either solid or void. The SIMP family turns this into a smooth problem by relaxing the design to a continuous density field in $[0, 1]$ and assigning an effective stiffness to each element through a material interpolation. Void is given a small ersatz stiffness so the finite element equations remain solvable. This relaxation makes the objective differentiable and allows efficient gradient evaluation for large models [8,10].

Penalization in SIMP is not cosmetic. It is the mechanism that steers the relaxed problem back toward black–white designs: weak penalization permits gray "mixtures," while stronger penalization biases the optimizer toward near 0/1 densities. Continuation (gradually tightening the interpolation) is therefore a principled pathway from an easier relaxed problem to a crisper design, rather than a tuning trick [8,10].

What SIMP guarantees and what it does not should be stated clearly. The state problem is well behaved, gradients are cheap and smooth, and multiple load cases fit naturally. However, SIMP alone does not ensure sharp boundaries or a physical feature size; without the regularization discussed previously, checkerboards and mesh patterns can appear. In practice, SIMP is paired with a fixed length scale (filters/projection or perimeter terms) to restore well-posedness and manufacturability [6,8,9,10].

For this study, SIMP provides a principled, reproducible baseline against which level-set can be compared. Both use the same physical length scale, so differences observed reflect formulation and solver rather than mesh artifacts. The smooth, relaxed structure of SIMP also justifies gradient-based optimizers (OC/MMA) in our setup [10,11].

2.5 Boundary conditions and their effect

Boundary conditions set how loads or heat enter and how motion is restrained. Small changes can lead to different topologies even when other settings are fixed.

Structural. Fixed, pinned, or symmetry supports remove selected motions. Very small support patches concentrate reactions and create thin members unless a length scale is enforced; larger pads or realistic fixture areas distribute loads and yield robust members. Point forces can cause mesh-dependent peaks and spindly links; distributed pressures or remote loads are more stable. Poor constraint sets leave near rigid modes that corrupt gradients or frequencies; tool defaults (weak springs, MPCs) can mask this and shift compliance across solvers [1,3,4,5,10]. Symmetry planes reduce search space but can exclude better asymmetric layouts when loads are not perfectly symmetric.

Thermal. Fixed temperatures, fluxes, and convection produce different layouts; small Dirichlet sinks favor trunk-like "heat pipes," while widespread convection favors fin arrays. Because the focus is structural manufacturability, these are treated as context and BC areas and coefficients are fixed precisely when used [2,5,10].

Manufacturability interaction. AM overhang limits and build direction bias members; supports near the base differ from side supports. Machining draw directions and minimum radii can conflict with tiny support footprints; larger pads often yield tool-accessible designs. Since BCs steer stress/heat paths, manufacturable members are directly affected by how process constraints align with those paths.

Cross-tool notes. Ansys and Comsol differ in default stabilization and constraint enforcement (weak springs, null-space handling, remote constraints). Two "identical" models can show measurable compliance or temperature shifts unless BCs match in area, type, and implementation. Accordingly, support regions are padded to a fixed area, point loads are avoided, distributed loads are preferred, and exact BC definitions are published to aid replication [1,3,4,5,10].

2.6 Knowledge gaps, inaccuracies, and contradictions

Many studies use older software; defaults and solvers change with release, making cross-paper comparisons difficult [1,3,4,5]. Key parameters (mesh size, p, filter radius, move limits, stop rules) are often underreported, hindering replication [1,3,4]. Manufacturability is sometimes applied only after the run, which hides trade-offs and can erase gains [1,4,6,7]. Experiments are limited and sometimes disagree with simulations, raising questions about settings and mesh dependence [1,4]. Cross-tool work often lists features rather than running end-to-end matched cases, so tool vs. setup effects remain unclear [5]. Level set and phase-field give smooth shapes but are not always matched to SIMP with the same length scale and constraints, weakening claims of advantage without a shared plan [2,10,12]. Recent reviews call for robust, physically based length scales and clearer reporting [9,10].

2.7 Theoretical background and workflow

A formulation (SIMP, level set, phase-field, or TD) is chosen, the governing physics and objective are set, regularization is applied to impose a physical length scale, manufacturing constraints are enforced during the run, gradient-based optimization with continuation is performed, and predefined metrics with convergence and timing are reported. This subsection states the sequence only; equations and constraint details appear above.

2.8 Thematic position of this work in the overall context

A controlled, cross-tool baseline is provided that isolates method and software effects by holding the problem definition fixed and enforcing manufacturability during optimization. The novelty lies in the fairness protocol and matched reporting across tools; high-level aims and scope are in the Introduction, and technical specifics are detailed earlier in this review.

2.9 Broadened software scope

Ansys advertises real-time topology and AM controls; in practice it is strong for structural SIMP with clear length scale and symmetry tools, but results depend on defaults so settings must be reported [1,3,4,5,9,10]. Comsol advertises multiphysics and level set; in practice it is good for coupled cases and sharp boundaries, while structural runs need careful filters [2,5]. Abaqus with Tosca supports topology/shape under nonlinearity and contact and is useful for draw/casting style rules, though setup is heavier [5]. Altair Inspire/OptiStruct emphasize manufacturability and patterns and are fast for SIMP concepting with minimum member size; detailed CAD rebuild is common [5]. nTop focuses on AM pipelines and lattices for post-processing/export rather than full solves [5].

3 Methodology

This section describes the numerical workflow used to generate all results. It also defines the exact modelling choices needed for a fair comparison between methods.

We proceed in two stages. First, we introduce a minimal pseudo 2-D benchmark plate that is easy to reproduce and sensitive enough to reveal differences in convergence, mass control, and stiffness. This benchmark is used to document our baseline settings, to verify that the chosen length scale produces stable layouts, and to provide iteration counts and final metrics on an identical case. Second, we apply the same protocol to a real suspension bracket design problem, where non-design regions are defined from assembly interfaces, three load cases are applied, and the resulting geometries are evaluated under the same reporting scheme. The case is pseudo 2-D, as it is a 3-D geometry, where the span of the geometry is one cell in the mesh. This decision was made as a controlled simplification as it made it significantly easier adapt the meshing protocols to the 3-D real case.

For each study, we report the optimisation method, iteration count, final mass or volume fraction, and the primary performance metric (strain energy or compliance). Where relevant, we also report stress summaries and provide qualitative observations on geometric quality and manufacturability.

3.1 2-D plate benchmark (reproducibility case)

- **Geometry:** 2000 mm $\times$ 1000 mm rectangle, left edge fully fixed.
- **Load:** distributed 1000 N load on a 100mm segment on top of the top right corner.
- **Material:** Structural Steel ($E = 210$ GPa, $\nu = 0.30$).
- **Objective:** minimum strain energy (compliance).
- **Volume constraint:** retain 50% of the initial mass.
- **Minimum member size:** $r_{\min} = 3$ mm (identical across all methods).

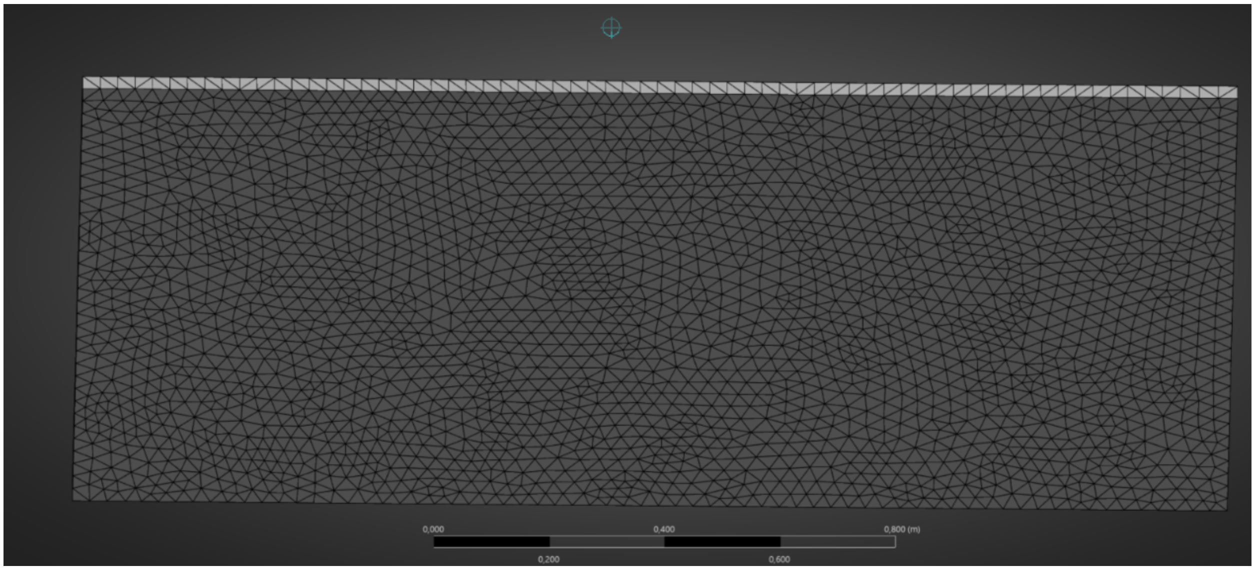

Figure 1: The pseudo 2-D nature of the geometry is visible in the mesh

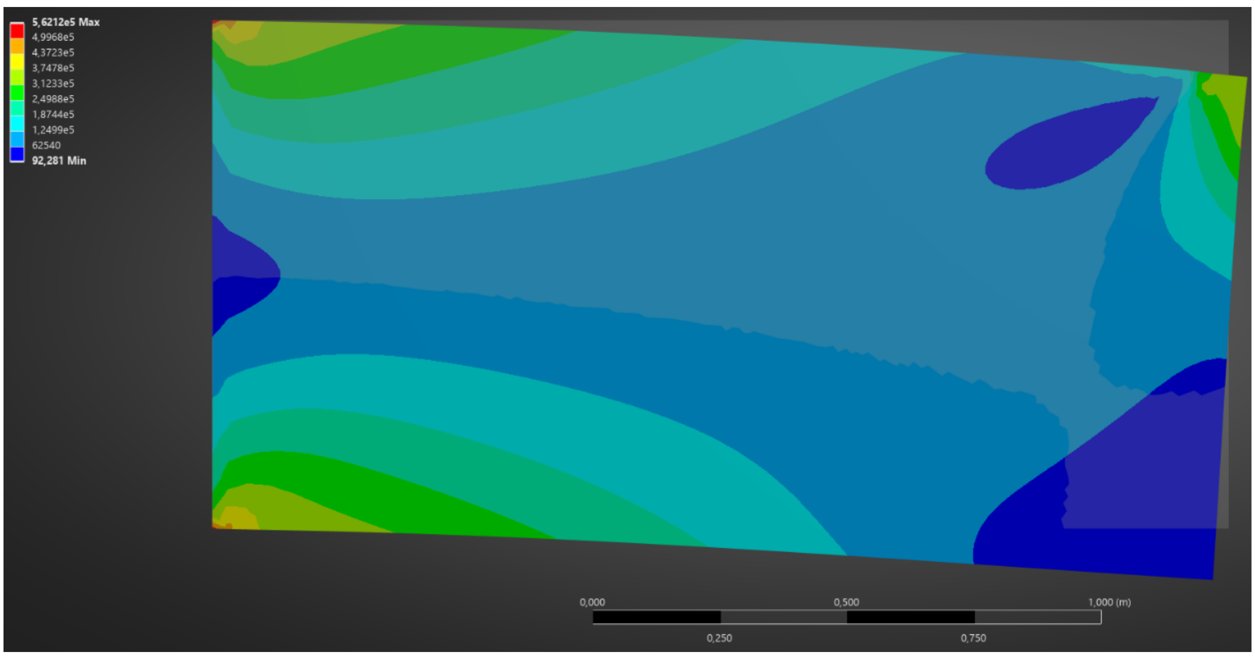

Figure 2: Benchmark plate: finite-element mesh (top) and stress field under the $1000\,$N load (bottom).

A minimum member size was enforced to keep the geometry pseudo 2-D, otherwise default settings were left unchanged.

The numerical outcomes (iterations, final mass and compliance) for this benchmark are collected in the following section.

3.2 Benchmark plate: final geometries

Figures below are the five final designs obtained with different topology-optimisation methods in Ansys Mechanical. All runs share exactly the same mesh, load, volume target and minimum member size defined in Section 3.1.

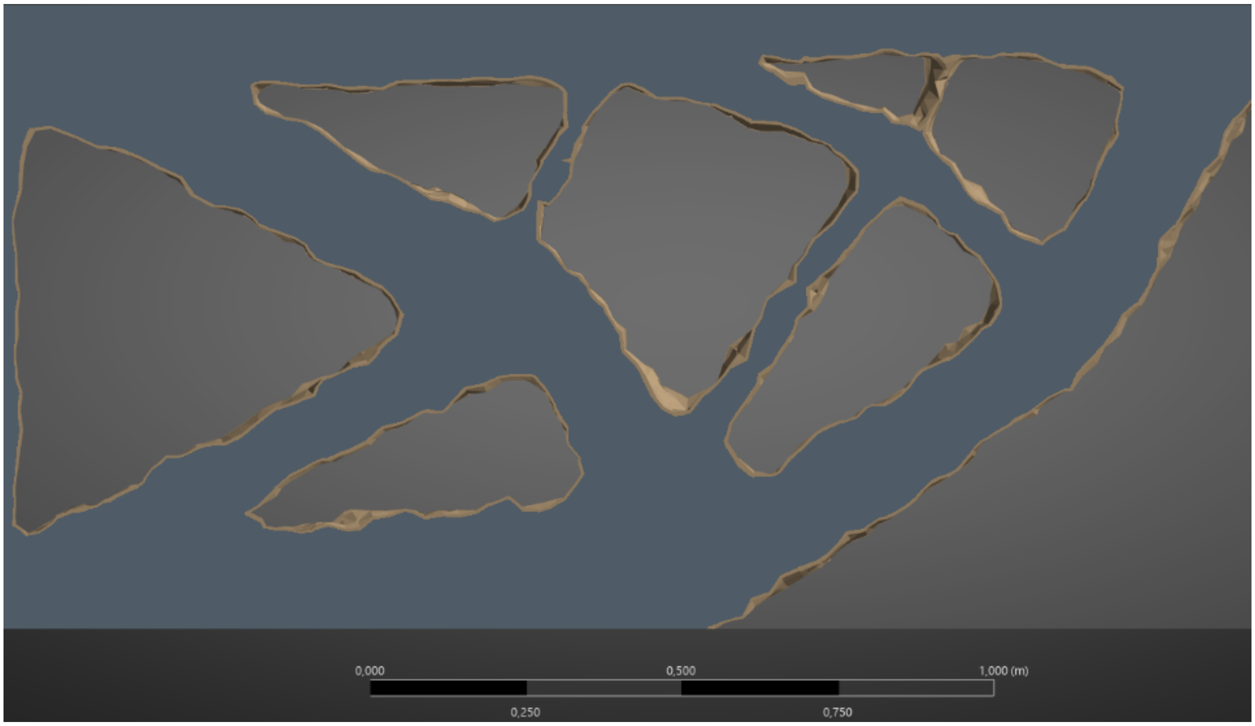

Figure 3: Optimised benchmark plate - SIMP with penalisation factor $p = 3$.

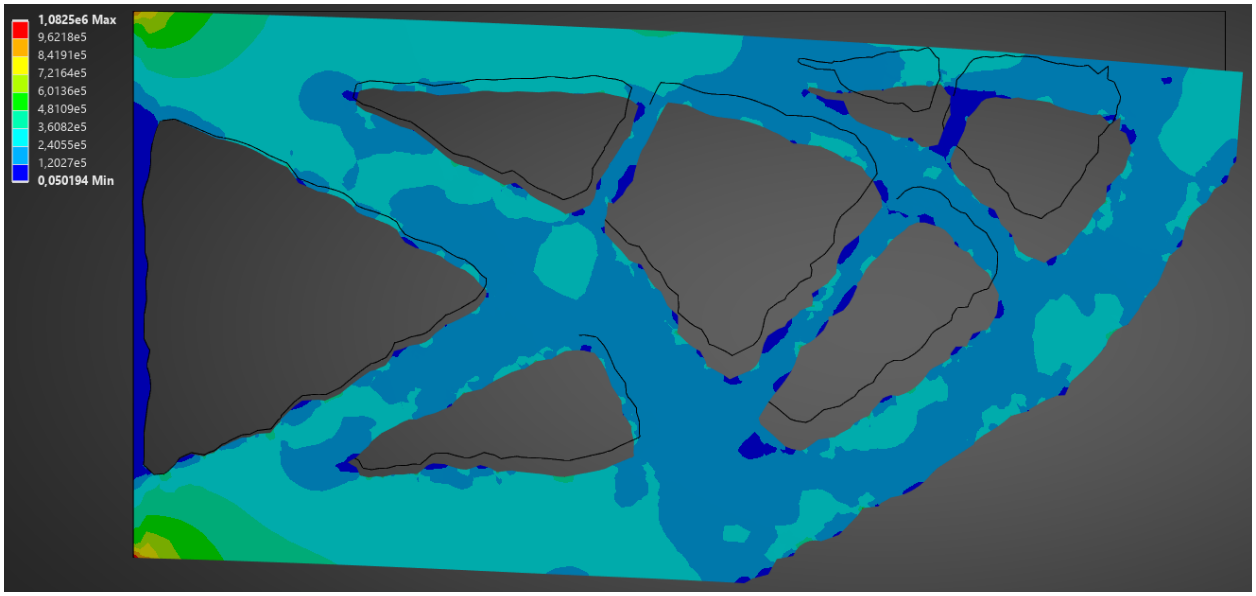

Figure 4: Equivalent (von Mises) Stress

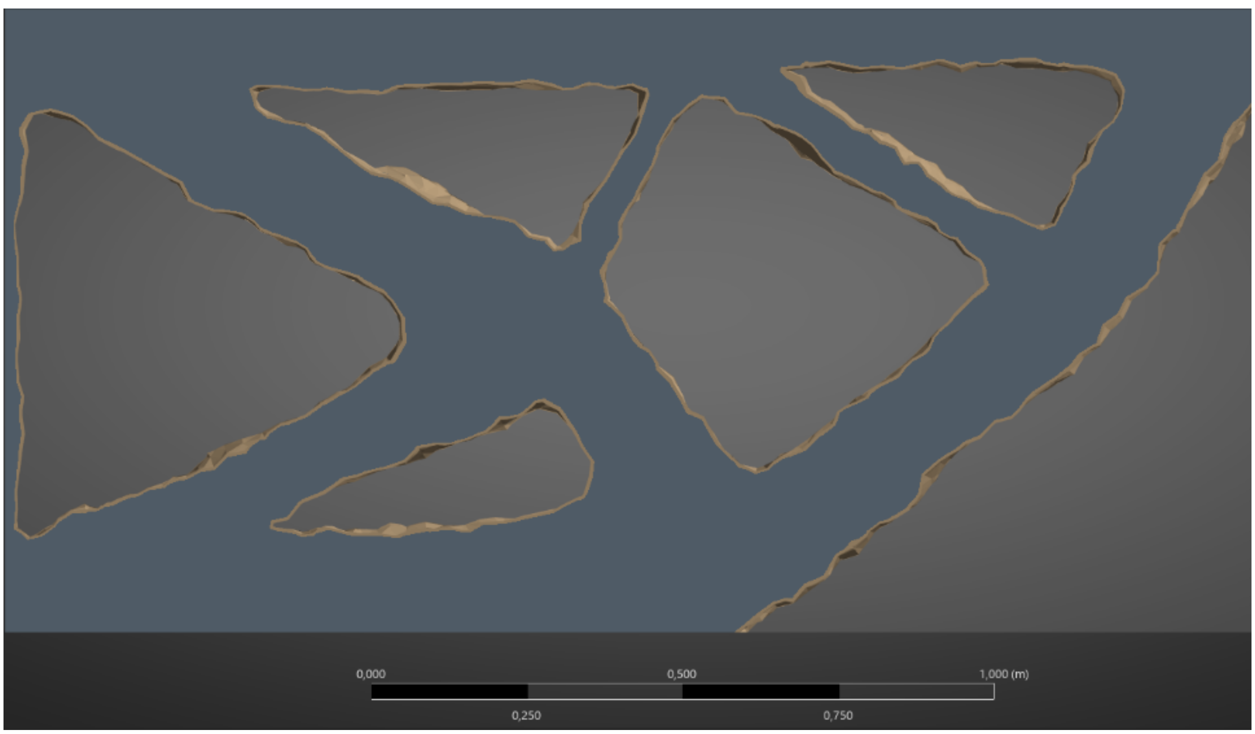

Figure 5: Optimised benchmark plate - SIMP with penalisation factor $p = 4$.

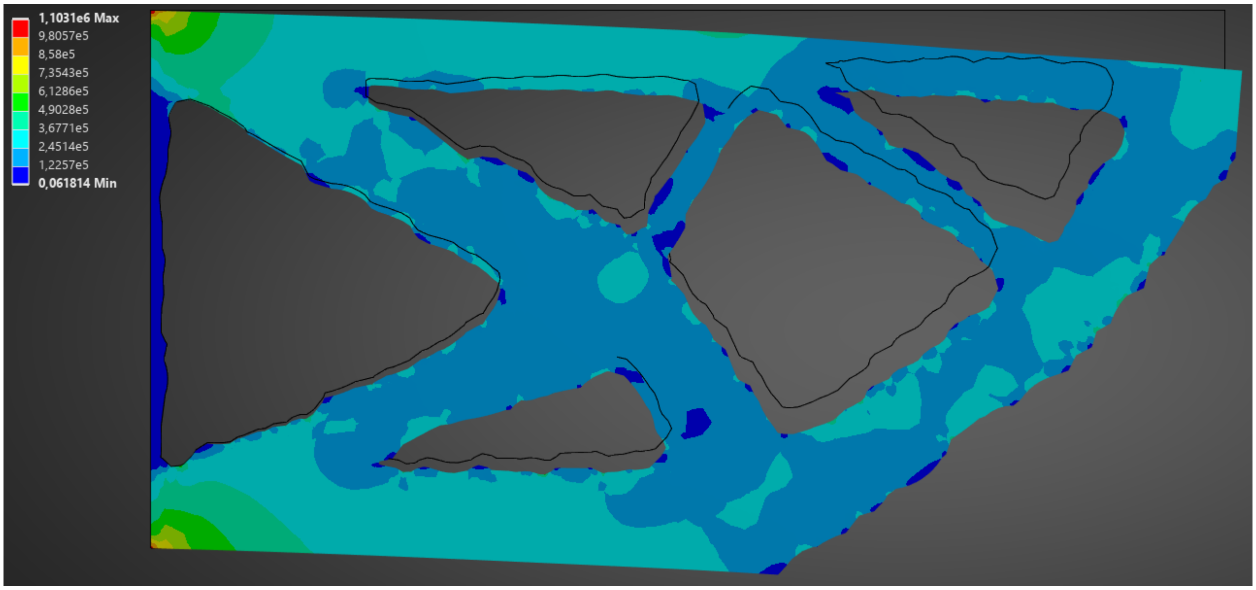

Figure 6: Equivalent (von Mises) Stress

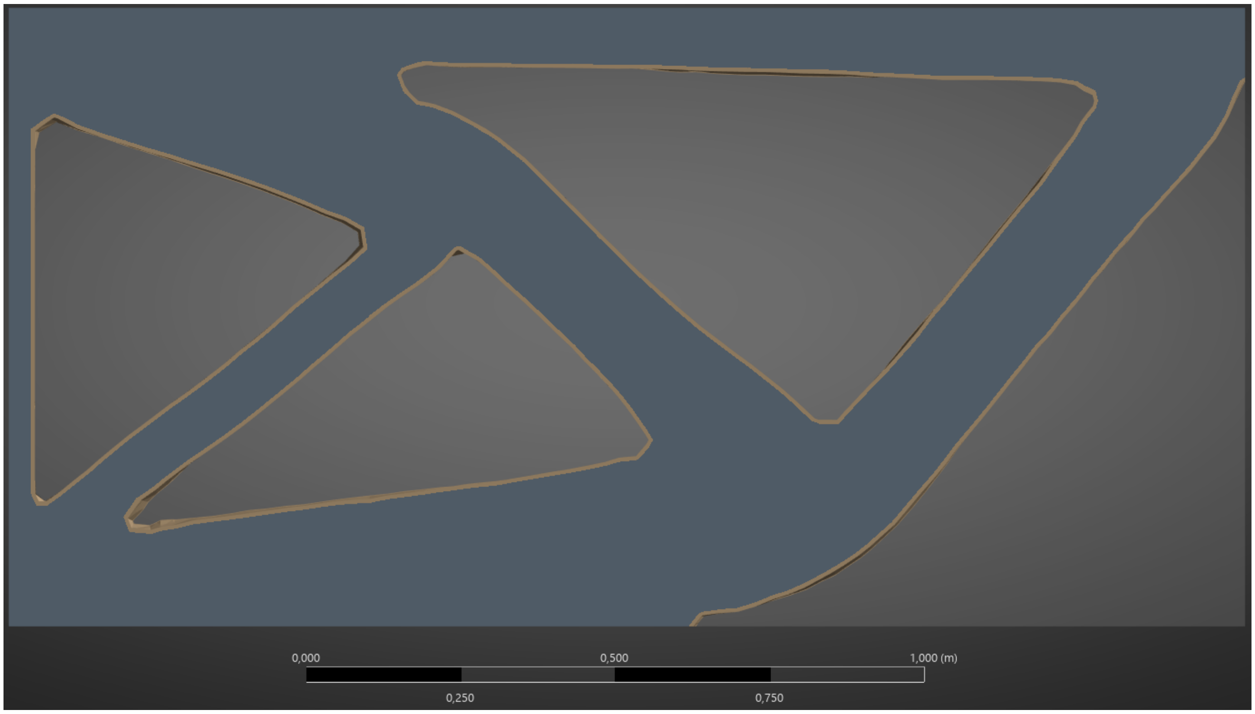

Figure 7: Optimised benchmark plate - level-set method.

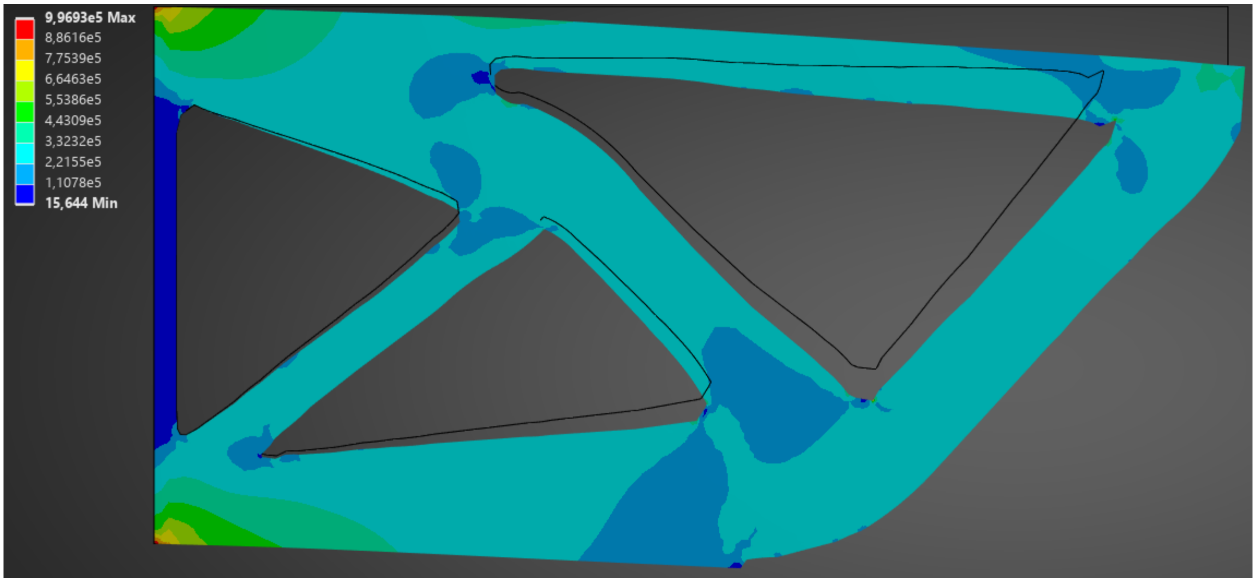

Figure 8: Equivalent (von Mises) Stress

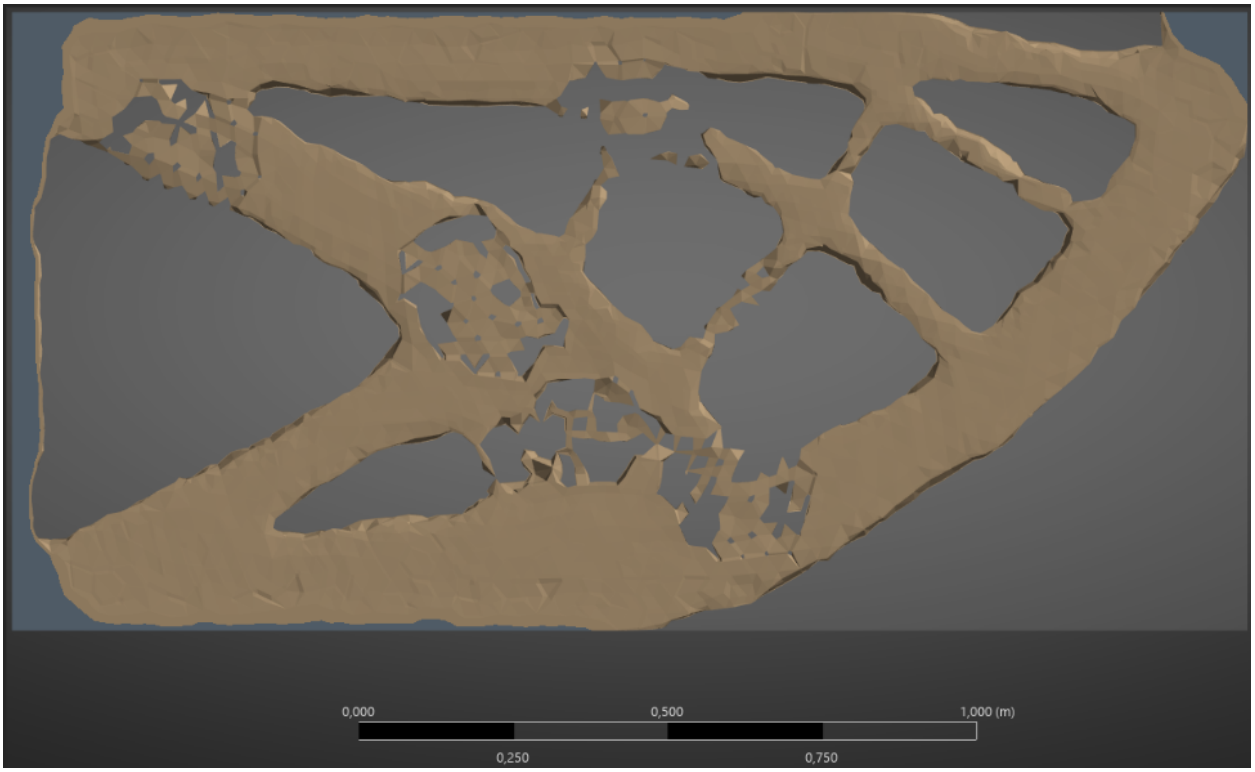

Figure 9: Optimised benchmark plate - mixable-density method.

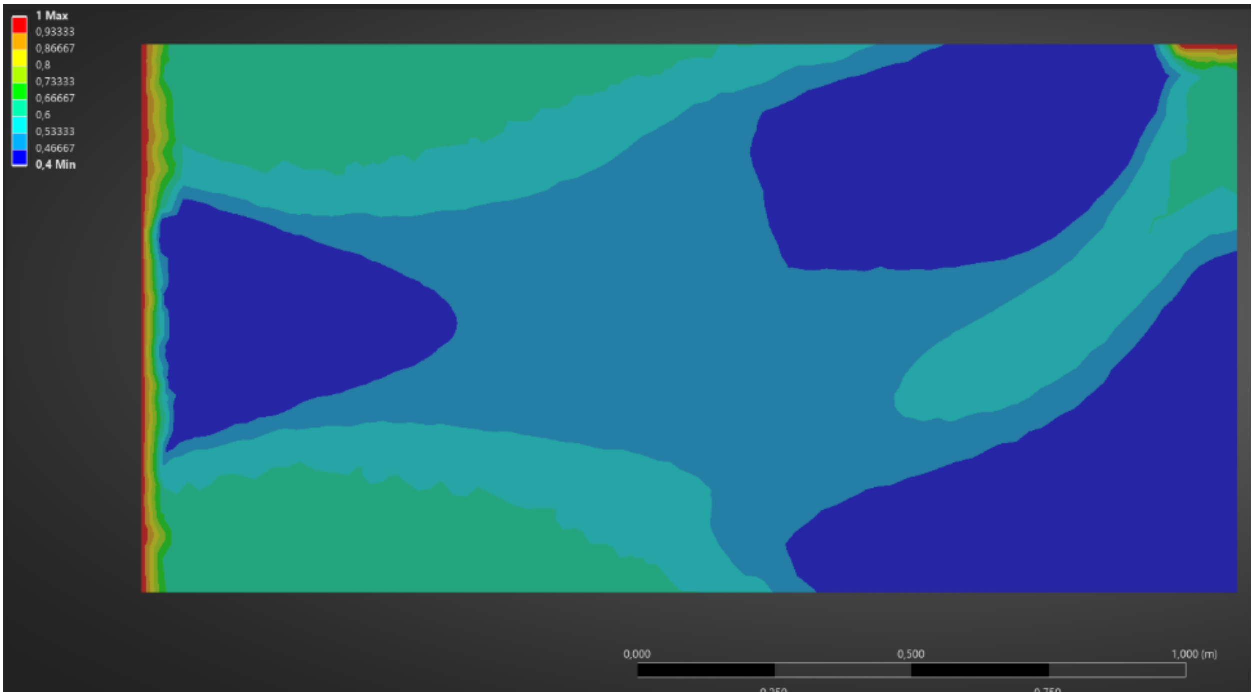

Figure 10: Optimised benchmark plate - cubic-lattice optimisation.

3.3 Key performance metrics

Method	Iter.	Mass. [%]	Final Mass	Strain Energy
SIMP ($p = 3$)	39	50.48	237.76	**4.993e-3 J**
SIMP ($p = 4$)	34	50.50	237.83	4.998e-3 J
Level-set	30	49.93	**235.15**	5.153e-3 J
Mixable density	41	52.32	246.43	n/a
Lattice (cubic)	12	50.42	237.46	n/a

Table 1: Iterations, final mass fraction and final mass for each method.

Observations.

- The *level-set* method produced the lightest design (about 1 % lower mass than the best SIMP run) while finishing in the fewest iterations of the full-density methods.
- *SIMP (p = 3)* produced the stiffest geometry, with the difference between penalization factors 3 and 4 almost negligible
- Raising the SIMP penalization from $p = 3$ to $p = 4$ cut five iterations but left the mass unchanged, indicating diminishing returns beyond a moderate penalization.

- The *Mixable-density* run overshot the volume target with the biggest margin, confirming the cost of excessive porosity.

- Different manufacturing constraints or filters would have to be studied and applied to make the *Mixable-density* geometry viable for practical use and to obtain the strain energy.

- The *Lattice* optimiser converged in only twelve iterations and matched SIMP mass; its graded infill may be attractive for additive manufacturing even if global stiffness is not improved.

3.4 Applying topology optimization in a real usecase

As a practical example, we will take a look at the process of designing a suspension bracket for a race car using topology optimization. Race car design is an excellent case study for topology optimization, as high stiffness at a minimal mass is highly desirable, and more costly manufacturing methods can commonly be used as the quantity of parts for a prototype like this is low very low. This allows for less constrained designs and more focus on performance. Suspension components in a race car are generally some of the most critical components on the car, regularly receiving the highest physical loads in a broad range of different load cases, and thus a structural suspension bracket from Metropolia Motorsports HPF025 car was chosen to optimize as a case study. The target of the design was to keep the same compliance as the original manually designed part at minimum, while reducing the mass of the part. A study on the original bracket was conducted using the original CAD files to determine the leading constraints for the part design, as well as the target performance metrics.

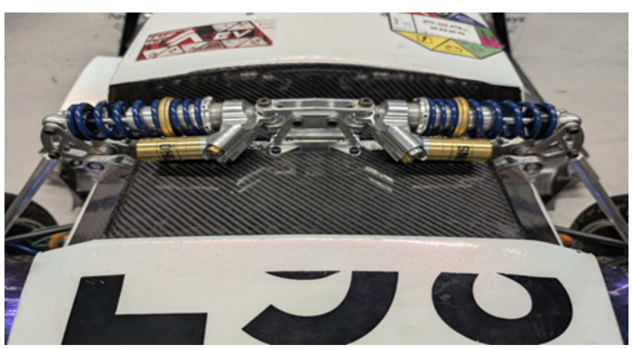

Figure 11: Inboard front suspension of the HPF025 race car with twin coilover dampers and rocker linkage; the bracket subject to optimization interfaces to the chassis in the middle via a four-bolt pattern.

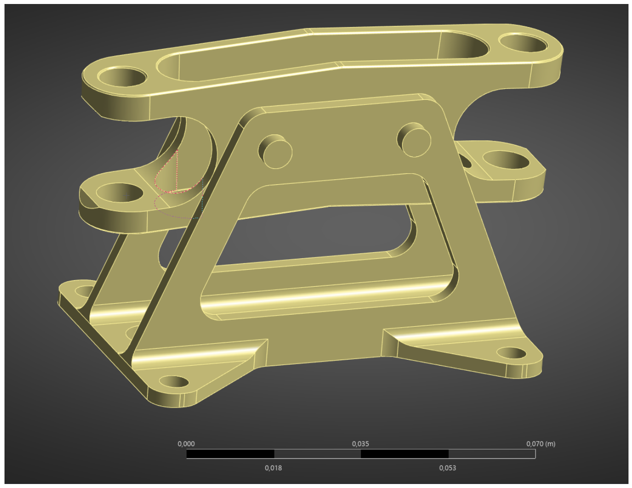

Figure 12: The CAD model representing the original, manually designed part.

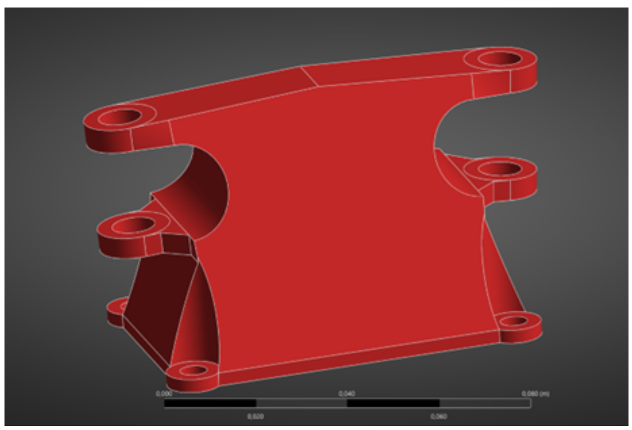

Figure 13: The optimization domain of the suspension bracket showing the four fixed bolt-hole locations, and the mounting points for the coilovers.

The bracket attaches to a rigid surface on the car with four bolts through the hole locations at the bottom of the part. The hole locations were predetermined by the previously used design, as it was determined critical that the suspension geometry itself remains without changes. The part only ever sees a load straight along the axis of the shocks, which exists in an angle scope very close to normal to the centerline through the holes where they attach. However, as the springs compress, the angle changes slightly, which means a bearing load can not be used as the load is not completely normal to the hole faces. What also has to be considered, is the load on the part is rarely symmetric as the car corners more than it drives in a straight line, and thus three different load cases have to be accounted for: asymmetric loading from both sides respectively during cornering, and a symmetric loading from braking in a straight line.

The material that will be used for the part is the aluminium based alloy AlSi10Mg, which is commonly used in aluminium 3D printing. The methods used for the real case, are Ansys' built-in level-set, and density-based algorithm. This is because based on the benchmark cases these methods seemed most suitable for the application. With data obtained from the car, the extremes of the three mentioned load cases were determined and a sufficient and a safety factor was applied. Important to note here is the car only weighs around 200kg.

3.5 Case Setup

The parts in the assembly will be bolted together, and as such, the surfaces where the bolts interface with the part were excluded from the optimization region.

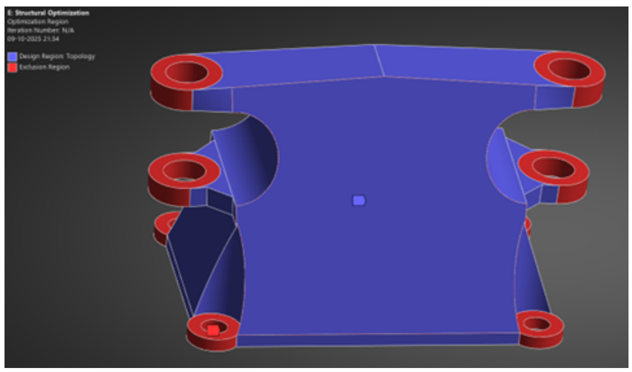

Figure 14: The region marked in blue is the optimisation domain; the segments marked in red are the excluded bolt areas.

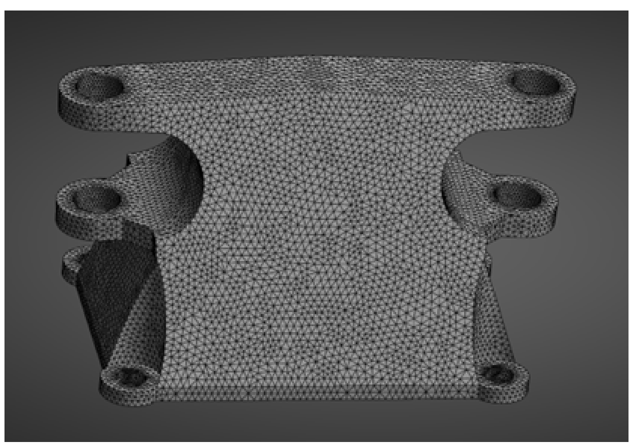

Figure 15: Finite-element mesh for the suspension bracket study.

The mesh was generated at the highest available preset resolution 7; preliminary tests showed that coarser meshes produced non-feasible geometries while computation time was still acceptable with the fine mesh. Further mesh studies with manual mesh settings could in the future be evaluated, however, as the results later will reveal, the mesh used here generated satisfactory results.

The three load cases are:

- **Symmetric braking load** - both dampers compress equally as the weight of the car shifts forwards in a straight manner.
- **Asymmetric cornering 1** - left damper compressed, right relaxed, as the weight of the car shifts to the left when turning right.
- **Asymmetric cornering 2** - right damper compressed, left relaxed, as the weight of the car shifts to the right when turning left.

The loads are applied as nodal forces along the damper axis, evenly distributed across both hole locations; fixing the four chassis bolt holes prevents rigid-body motion. Important to note: the bolt holes in the load cases deform in an unnatural manner as nodal forces had to be used in place of bearing loads. However, this has no significant effect on the optimized geometries, as the unnatural deformation exists within the regions excluded for optimization Figures 16-18 illustrate the three cases.

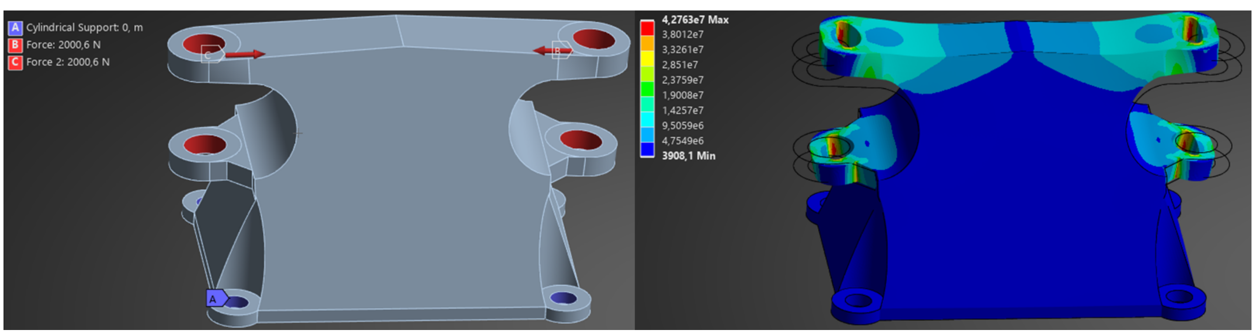

Figure 16: Symmetric braking load: boundary conditions and von Mises stress.

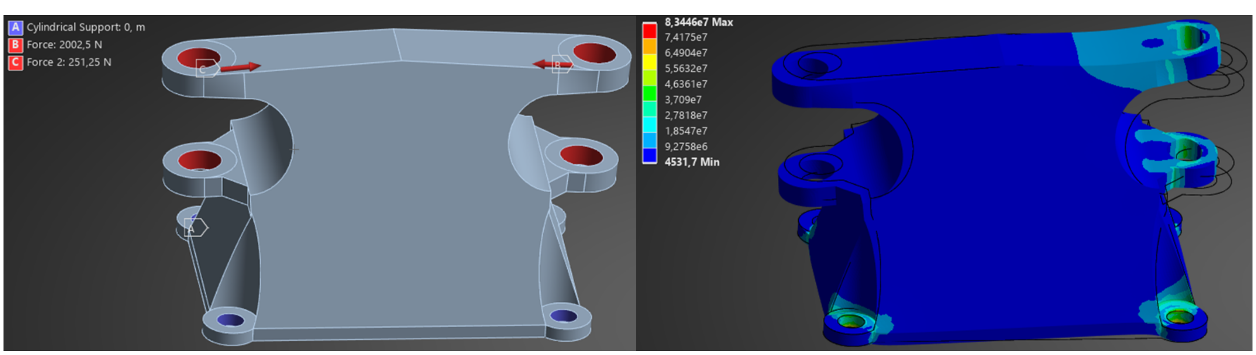

Figure 17: Cornering load 1: left-side compression: boundary conditions and von Mises stress.

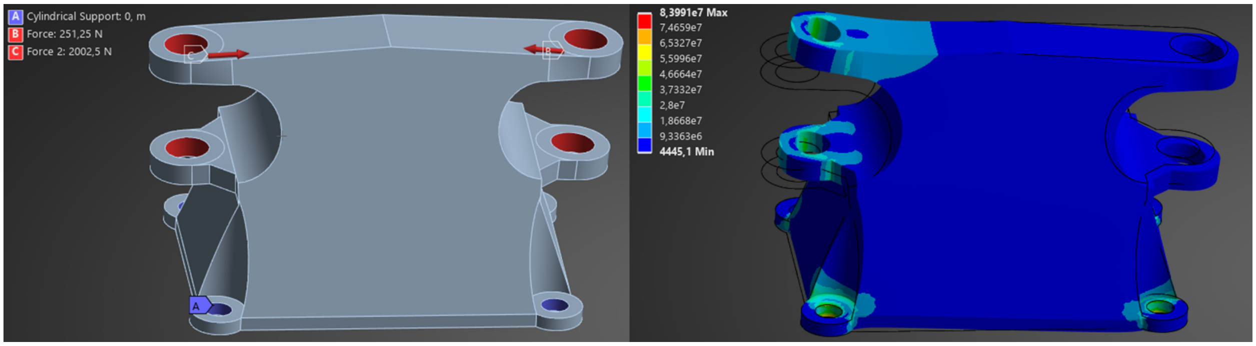

Figure 18: Cornering load 2: right-side compression: boundary conditions and von Mises stress.

The structural analyses above feed directly into the Ansys structural optimisation module. The response constraint is set at 20% mass fraction of the optimization region.

4 Results

4.1 Original bracket results under the same load cases

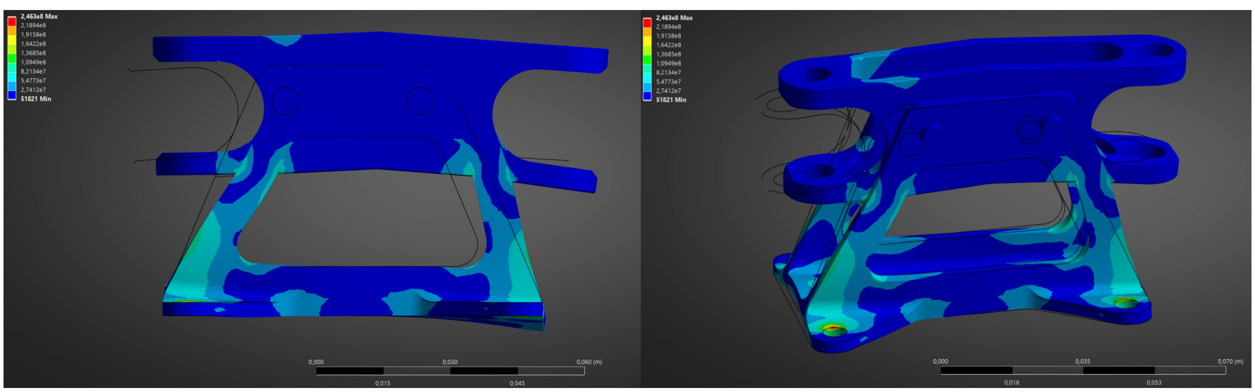

Figure 19: Original bracket, asymmetric cornering load, von Mises stress, two views.

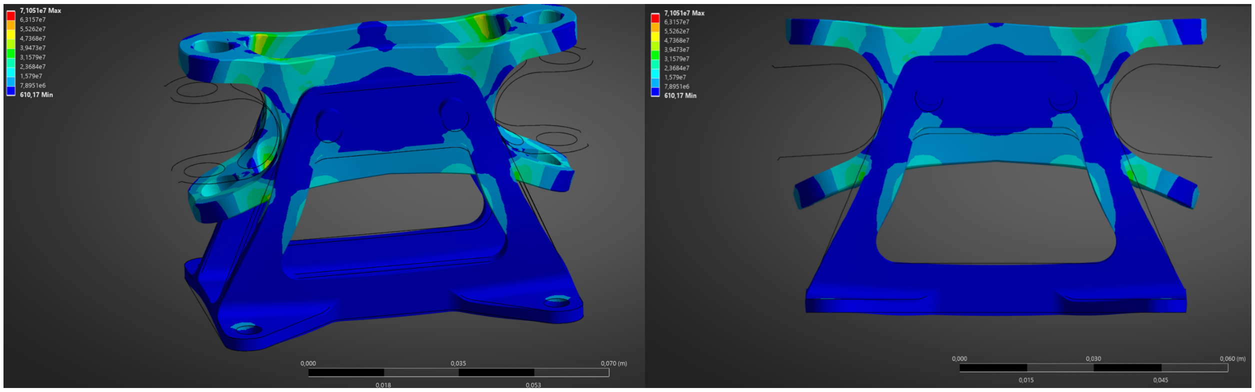

Figure 20: Original bracket, symmetric braking load, von Mises stress, two views.

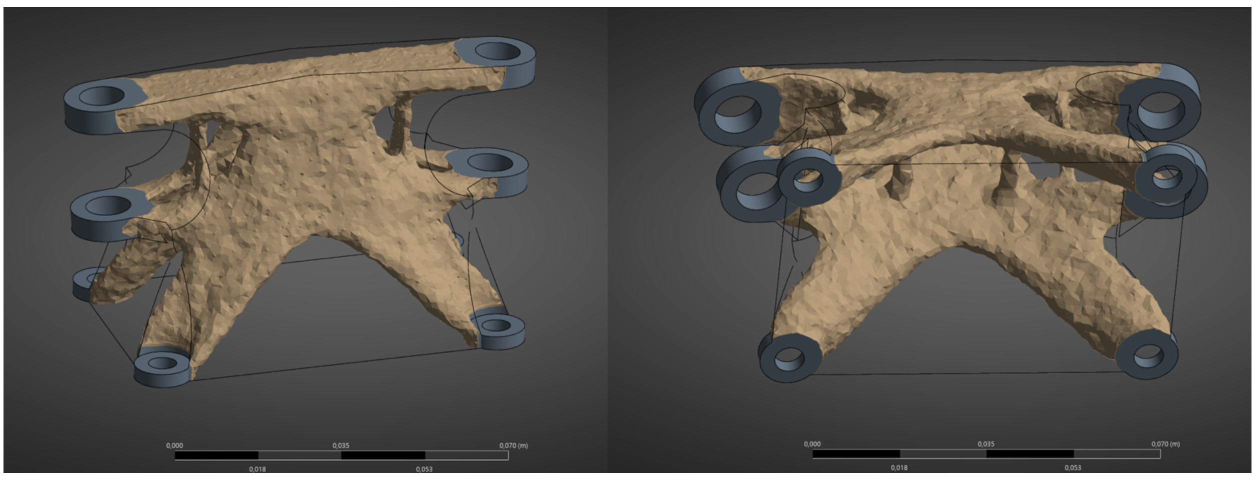

Figure 21: Bracket geometry 1: The geometry resulting from using SIMP with $p = 3$.

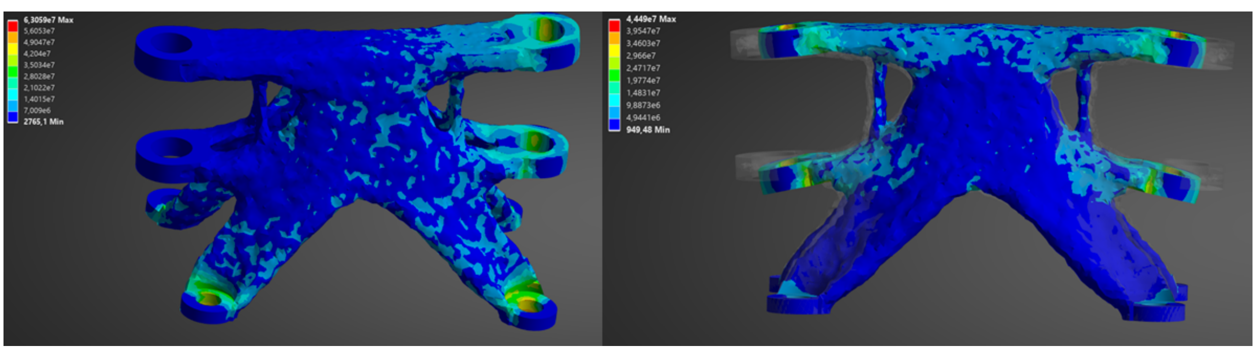

Figure 22: Bracket geometry 1: von Mises Stresses

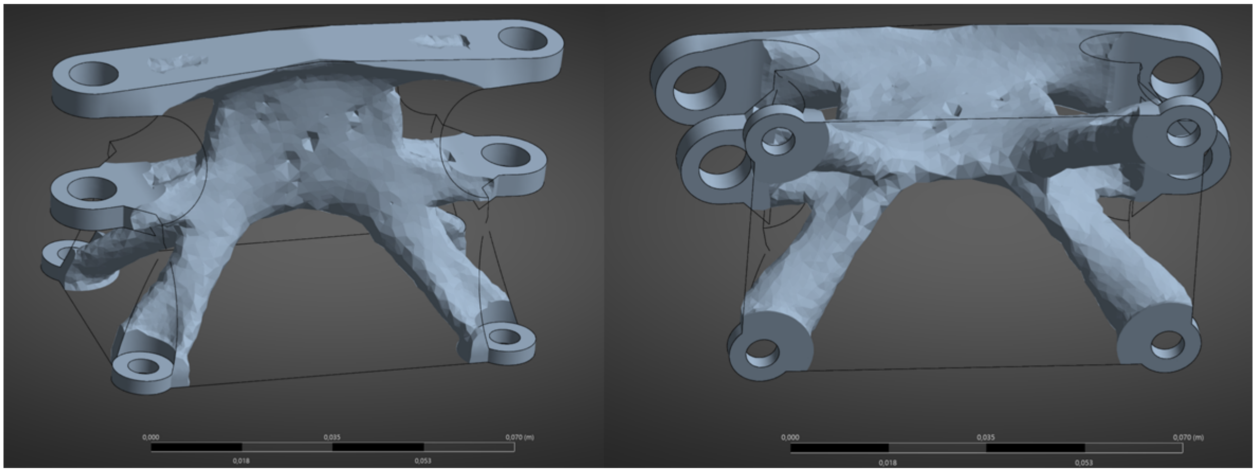

Figure 23: Bracket geometry 2: The geometry resulting from Ansys built in level-set method.

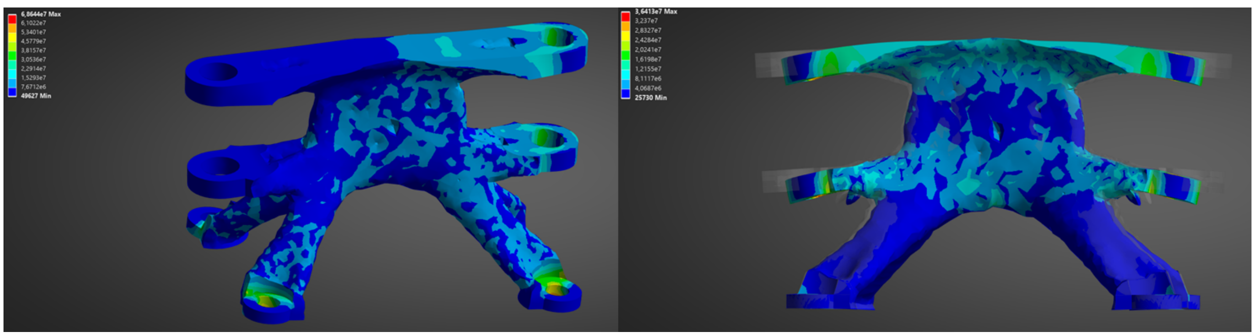

Figure 24: Bracket geometry 2: von Mises Stresses

4.2 Key performance metrics

Method	Iter.	Final Mass [g]	Mass [%]	Strain Energy [J]	Average von Mises Stress [Pa]
Original	n/a	288.1	100.0	4.961e-2	1.567e7
SIMP ($p = 3$)	**22**	230.1	79.9	**9.561e-3**	**9.255e6**
Level-Set	28	**213.2**	**74.0**	1.139e-2	1.026e7

Table 2: Comparison of the performance metrics of the different parts.

Observations:

- Both methods achieved the goal of reduced weight while at minimum maintaining stiffness higher than that of the original part

- *SIMP* produced the stiffest part at around 16% stiffer than than the level-set method and around 80% stiffer than the original part, while doing so in fewer iterations than the *level-set* method.
- *Level-set* Produced the lightest design, just like in the benchmark case, this time around 7% lighter than SIMP, and around 26% lighter than the original part.
- The lowest average stress was achieved by *SIMP*, with an average stress around 10% lower than the level-set geometry, and around 40% lower than the original geometry.
- At a glance, both methods produced what look like feasible geometries, and no obvious artifacts can be observed.
- The two generated parts differ significantly in geometry even if some observable trends also emerge.
- While different filters and constraints could be applied to change the final part, no filters or constraints beyond the defaults in Ansys were applied to make the comparison as neutral as possible.

With those initial observations, we can take a deeper look at what the results represent. Starting with the similarities, It is clear that the ”legs” of the parts look quite similar, and the overall distribution of material is trending in the same direction. Both parts generally lack small geometries or signs of significant checker-boarding. Neither method produced any significant discontinuities or strange voids. Both methods out-perform the original part in all categories with a comfortable margin, meaning that there is still some headroom for further optimization and/or safety factors.

Observing the many similarities, quite a few notable differences also stand out. While the part produced by the level-set method is the lightest part of the three, we can also clearly see that compared to the slightly heavier part made by SIMP, it does sacrifice stiffness and strength significantly to achieve the lower weight. The part generated with SIMP also has a large hollow mid-section, while the level-set generated a plate-like structure in the same place. Outside of the designed loads this will make for a big difference.

Overall in these metrics SIMP seems to have a slight edge over the level-set method, however it is not possible to isolate if this is definite or if one method is significantly more limited by the lack of filtering or constraints. For that a study on how different filters and constraints affect the results would have to be conducted.

5 Discussion

Observing the results from both the benchmark case and the real case, clear and reproducible differences between the methodologies emerge. The benchmark case served as a test to decide upon which methods to use and compare in the computationally heavier real case, as well as to compare basic metrics of the methods themselves. As all controllable parameters were fixed between the methods, the differences in the outputs can be characterized by the methods used. As such it is easy to evaluate and isolate variables.

For the benchmark case, SIMP produced the stiffest results, with penalty factor having negligible effect on the stiffness even if the geometry generated differed significantly. This indicates that penalization beyond moderate levels provide diminishing returns in terms of performance, but might be a useful tool when it considering manufacturability or part feasability. The level-set method on the other hand provided the lightest geometry, providing sharp and smooth boundaries but at a small cost in stiffness. Within the scope of this study the mixable density method struggled to produce feasible results, as major checker boarding and porosity was prominent in the final geometry. This was not however without promising potential. As filters or constraints were omitted for a fair comparison, mixable density might only show its true potential with some additional tweaking. Mixable density also opens the potential for multimaterial properties which were not explored in this report.

The lattice-based optimizer converged rapidly and achieved target mass but is not directly comparable in stiffness because its effective properties depend strongly on micro-architecture rather than global topology. As such it was mostly included as context to highlight the possibility of lattice-based optimization.

In general these findings align with reported literature, as SIMP tends to produce robust load paths with strong mass-stiffness performance, while level-set methods generally tend towards smoother and sharper boundaries. Level-set also showed a tendency to undershoot the mass response target in contrast to SIMP which generally ended up slightly above.

Moving on to the suspension bracket subsequently studied, two targets were clearly outlined. The goal was to use topology optimization to replace a manually designed part with a generated design improving both on mass and stiffness. From the benchmark study, SIMP and level-set method was selected as methods to evaluate for this purpose, as they showed least compromise for the task. As the mass response constraint of 20% easily outperformed the original part in all metrics, the weight reduction could possibly be pushed further without significant structural compromise, and as filtering and constraints were kept at a minimum, further performance could possibly be found through those means in future studies. However proven here is that the targets were well within reach and topology optimization provides significant benefit over manual design.

As the geometries observably provided significantly different resulting geometries, this also opens opportunities for further study, where factors like fatigue and vibration resistance and more could provide significant differences in long-term performance. For a race-car this would be highly relevant even if beyond the scope of this study. The mixable density also represents a method that could be studied more in depth for the purpose at hand, as this study likely only scratched the surface of that method.

6 Conclusions

This paper's main contribution is to connect the reported Ansys outcomes to the standard theory used in topology optimisation. For linear elasticity under load control, equilibrium $Ku = f$ implies that compliance $C = f^T u$ and strain energy $U = \frac{1}{2}u^T Ku$ are equivalent up to a constant factor, $C = 2U$. Therefore, lower strain energy directly indicates higher stiffness for the same load definition and admissible volume fraction, which justifies using strain energy or compliance as the primary structural objective.

A second conclusion is that length-scale control is not optional but required for well-posedness and meaningful comparison. The literature shows that unregularised compliance minimization is ill posed and mesh dependent. Practical methods restore stability by imposing a fixed physical length scale, typically via filtering and projection in density-based SIMP, or via curvature and smoothing controls in boundary-based level-set formulations. Under this shared requirement, differences between SIMP and level-set are interpreted as differences in representation and regularization: SIMP optimizes over relaxed densities with penalization and continuation toward near-binary designs, while level-set evolves an explicit interface via shape sensitivities, often producing sharper, CAD-ready boundaries.

Thirdly, with a practical case presented in the report, the scope is broadened and through applying the theory to a tangible use case, a more intuitive understanding of the topic is developed. With the targets set out early, the process of generating a part that satisfies the targets is described. Since the nature of the study means some factors have to be omitted, those are subsequently discussed. Further having two different cases where the same methods are compared also gives a clearer picture of what characteristics discovered can be linked to the method, as they are repeated multiple times. This means a more objective overview of the strengths of SIMP and level-set method respectively is formed alongside the practical exploration of topology optimized design.

Finally, the controlled benchmarking framework itself is a meaningful outcome of the work. The literature repeatedly notes that reported differences across studies often stem from inconsistent meshes, solver defaults, undocumented filters, or shifting manufacturing constraints rather than actual methodological differences. By fixing all modelling and physical parameters and enforcing manufacturability constraints during the optimisation rather than after it, this study demonstrates a replicable path for cross-method comparisons. Future work can build on this framework by broadening the range of constraints, extending to nonlinear or multi-physics scenarios, or comparing a wider set of commercial tools under the same fairness protocol. In this sense, the presented results contribute both specific technical insights and a general methodology for more transparent and comparable topology-optimisation research.

6.1 Acknowledgements

We thank our advisor Rauf Talha for guidance throughout the project. We also acknowledge *Aalto University* for the course framework, and *Metropolia Motorsports* for the suspension bracket context and motivation for the real-case study.

7 References

1. Gunwant, Misra. *Topology Optimization of Sheet Metal Brackets Using ANSYS*. 2012 MIT International Journal of Mechanical Engineering, Vol 2 No 2, pp. 120–126)

2. Zhang, Li, Zhong, Xiang. *Structural Topology Optimization Based on the Level Set Method Using COMSOL* . *CMES* 101(1), 2014.

3. Bodog, Grebenian. *An integrated approach of topological optimization with parametric validation and response surfaces, using ANSYS*. *IOP Conf. Ser.: MSE* 1256, 012014, 2022.

4. Coşkun, Aydın, Koçak. *Topology Optimization of Snake-like Robot Limb Using Ansys*. *Journal of Materials and Electronic Devices* 5(1), 2023.

5. Struž, Hruzik, Klapetek. *Comparative Analysis of Different Softwares in Terms of Parameters Optimized by Topological Optimization*. *MM Science Journal*, March 2023.

6. Guest, James K., Jean H. Prévost, and Ted Belytschko. 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, 2004.

7. Langelaar. *An additive manufacturing filter for topology optimization of print-ready designs*. *Struct. Multidiscip. Optim.* 55, 2017.

8. Bendsoe, Sigmund. *Topology Optimization, Theory, Methods, and Applications*. Springer, 2003 and later printings.

9. Andreassen et al. *Efficient topology optimization in MATLAB using 88 lines of code*. 2011.

10. Sigmund, Maute. *Topology optimization approaches, A comparative review*. *Struct. Multidiscip. Optim.* 2013.

11. Svanberg. *The method of moving asymptotes, a new method for structural optimization*. *Int. J. Numer. Methods Eng.* 24(2), 1987.

12. Allaire, Jouve, Toader. *Structural optimization using sensitivity analysis and a level set method*. *J. Comput. Phys.* 194(1), 2004.