Proceedings of the Seventh International Conference on
Artificial Intelligence, Soft Computing, Machine Learning and Optimization,
in Civil, Structural and Environmental Engineering
Edited by: P. Iványi, J. Kruis and B.H.V. Topping
Civil-Comp Conferences, Volume 11, Paper 5.1
Civil-Comp Press, Edinburgh, United Kingdom, 2025
ISSN: 2753-3239, doi: 10.4203/ccc.11.5.1

HiCon-FEM: A Hierarchical Condensation Framework for Accelerated Topology Optimization

V. Yanes[1], N.-H. Kim[2] and F. J. Montans[1]

[1] Escuela Técnica Superior de Ingeniería Aeronáutica y del Espacio, Universidad Politécnica de Madrid, Spain
[2] Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, United States

Abstract

This work introduces a finite element framework to accelerate density-based topology optimization through reduced hierarchical bases and condensation. The mesh has microelements embeded into macroelements. The method projects fine-scale displacements onto a reduced basis that consists of boundary and internal modes, and performs the elimination of internal degrees of freedom at the macroelement level prior to global assembly. This results in a significantly smaller linear system while preserving compatibility with the standard SIMP optimization loop. The method requires no modifications to the optimization algorithm and remains robust across different filtering strategies. Through classical benchmark examples, the method demonstrates its ability to deliver high-resolution, mechanically accurate designs with reduced computational effort, making it a practical choice for large-scale topology optimization problems.

Keywords: topology optimization, finite element method, hierarchical basis, model order reduction, structural design, topology optimization, static condensation

1 Introduction

Topology optimization is key in modern component design, enabling the automatic generation of high-performance layouts for lightweight and mechanically efficient components. Since its inception through homogenization-based methods [1] and the development of the SIMP (Solid Isotropic Material with Penalization) approach [2,3], topology optimization (TO) has found many applications in aerospace [4], automotive [5], and civil engineering [6]. However, TO needs the repetitive solution of large finite element models, so the computational demands of solving the large-scale finite element systems at every iteration of the TO algorithm remain a critical bottleneck—particularly in high-resolution or three-dimensional problems.

Efficient solution of large finite element models is common to other problems, like multiscale ones. Many acceleration strategies have been proposed to address this challenge, including multigrid solvers (teamed with iterative methods) [7], multiscale frameworks [8], substructuring-based reduction [9], and neural network surrogates [10]. More recently, methods combining static condensation and model order reduction (MOR) have shown promise in reducing analysis time without sacrificing fidelity [11–13]. These methods can also be combined with the previous ones. While these efforts have improved performance in specific settings, few offer a general, scalable solution compatible with standard density-based optimization loops.

In this context, we propose a finite element framework that incorporates hierarchical condensation to accelerate structural analysis within topology optimization (named HiCon-FEM). The method reduces computational complexity by projecting element-level displacements onto a compact basis—comprising boundary and internal modes—and eliminating internal degrees of freedom prior to global assembly. Unlike traditional multiscale methods, HiCon-FEM requires no assumption of scale separation and integrates seamlessly into the classical SIMP loop.

The key ideas behind HiCon-FEM are presented herein. Its efficiency and accuracy are demonstrated through classical benchmark examples. The results indicate that HiCon-FEM can substantially reduce computational effort while delivering high-resolution, physically meaningful designs.

In the next sections, we first introduce the formulation of the proposed method, including its integration into the SIMP framework. We then validate its performance through numerical experiments on standard topology optimization benchmarks. Finally, we discuss the scope, limitations, and potential directions of the method for future work.

2 Methodology

HiCon-FEM accelerates topology optimization by reducing the computational cost of finite element analysis through a two-stage hierarchical condensation process. It is designed to integrate directly into the standard SIMP framework without requiring changes to the optimization algorithm. In this section, we present the methodology of the method as well as its integration with the SIMP algorithm.

2.1 SIMP-Based Topology Optimization

Although the ideas can be used in any TO procedure, because of their wide use, we adopt the standard SIMP formulation. In SIMP, the design domain is discretized into finite elements, each associated with a density variable $\rho_e \in [\rho_{\min}, 1]$. The objective is to minimize structural compliance under a material volume constraint:

$$\min_{\rho} \quad c(\rho) = \mathbf{f}^T \mathbf{u} \quad \text{s.t.} \quad \sum_{e=1}^{N_e} v_e \rho_e \leq V, \quad \mathbf{K}(\rho)\mathbf{u} = \mathbf{f}, \tag{1}$$

where $\mathbf{K}(\rho) = \sum_e \rho_e^p \mathbf{K}_0^e$ is the global stiffness matrix assembled from element-wise contributions. The penalization factor $p \geq 3$ promotes discrete designs, while a small lower bound $0 < \rho_{\min} \ll 1$ prevents singularities in regions approaching void.

The compliance sensitivity is obtained via the adjoint method:

$$\frac{\partial c}{\partial \rho_e} = -p\rho_e^{p-1} \mathbf{u}_e^T \mathbf{K}_0^e \mathbf{u}_e, \tag{2}$$

where $\mathbf{u}_e$ is the displacement vector restricted to element e. These sensitivities are used to update the design using the optimality criteria (OC) method, which remains one of the most effective approaches for density-based optimization problems.

To suppress numerical instabilities such as checkerboarding and mesh dependency, we apply filtering strategies based on the example. For the MBB beam, we use a classical sensitivity filter:

$$\widehat{\frac{\partial c}{\partial \rho_e}} = \frac{1}{\rho_e \sum_{j \in N_e} w_{ej}} \sum_{j \in N_e} w_{ej} \rho_j \frac{\partial c}{\partial \rho_j}, \tag{3}$$

where $w_{ej} = \max(0, r_{\text{min}} - \text{dist}(e, j))$, and N_e is the neighborhood of element e. For larger-scale problems, such as bridge and cantilever examples, we adopt a convolution-based density filter. A normalized Gaussian kernel centered at (i_0, j_0) with a standard deviation σ is defined as:

$$G(i, j) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(i - i_0)^2 + (j - j_0)^2}{2\sigma^2}\right), \tag{4}$$

and the filtered sensitivity is computed as:

$$\frac{\widehat{\partial c}}{\partial \rho_e} = \frac{G * \left(\rho_e \cdot \frac{\partial c}{\partial \rho_e}\right)}{G * \rho_e + \varepsilon}, \tag{5}$$

where $*$ denotes full convolution and ε is a small number to prevent division by zero.

These filtering techniques improve convergence and promote manufacturable black-and-white topologies. With the optimization framework established, we next introduce the HiCon-FEM condensation procedure that accelerates the analysis phase without altering the SIMP optimization loop.

2.2 Hierarchical Condensation Framework

The core idea behind HiCon-FEM is to accelerate the finite element analysis step by applying model order reduction at the element level before global assembly. The method relies on two consecutive operations: projection onto a hierarchical basis and static condensation of internal modes. These steps are performed independently within each macroelement, allowing efficient parallel implementation.

Each macroelement is composed of a structured grid of standard micro finite elements. Let $\mathbf{u}_e \in \mathbb{R}^{n_e}$ denote the vector of all displacement degrees of freedom (DOFs) associated with the microelements inside a macroelement. Instead of assembling $\mathbf{u}_e$ directly into the global system, we approximate it using a reduced vector $\mathbf{v}_e \in \mathbb{R}^{m_e}$ via:

$$\mathbf{u}_e \approx \mathbf{A}_e \mathbf{v}_e, \tag{6}$$

where $\mathbf{A}_e \in \mathbb{R}^{n_e \times m_e}$ is an interpolation matrix constructed from hierarchical basis functions. These include serendipity-type functions defined on the boundary and a set of internal "bubble" modes with vanishing trace on the boundary.

By substituting into the local equilibrium equation, we obtain a reduced stiffness system:

$$\tilde{\mathbf{K}}_e \mathbf{v}_e = \tilde{\mathbf{f}}_e, \quad \text{with} \quad \tilde{\mathbf{K}}_e = \mathbf{A}_e^T \mathbf{K}_e \mathbf{A}_e, \quad \tilde{\mathbf{f}}_e = \mathbf{A}_e^T \mathbf{f}_e, \tag{7}$$

where $\mathbf{K}_e$ is the standard elemental stiffness matrix assembled over all microelements within the macroelement, and $\mathbf{f}_e$ is the corresponding force vector.

To further reduce system size, we partition $\mathbf{v}_e$ into external (boundary) DOFs $\mathbf{v}_e^E$ and internal (bubble) DOFs $\mathbf{v}_e^I$. Assuming no external forces act on the internal modes, static condensation yields:

$$\mathbf{v}_e^I = -(\tilde{\mathbf{K}}_e^{II})^{-1} \tilde{\mathbf{K}}_e^{IE} \mathbf{v}_e^E, \tag{8}$$

which leads to a final condensed stiffness matrix defined only on boundary DOFs:

$$\hat{\mathbf{K}}_e = \tilde{\mathbf{K}}_e^{EE} - \tilde{\mathbf{K}}_e^{EI}(\tilde{\mathbf{K}}_e^{II})^{-1}\tilde{\mathbf{K}}_e^{IE}. \tag{9}$$

This condensed matrix $\hat{\mathbf{K}}_e$ is then assembled into the global system. Although the resulting matrix is generally denser, the number of global DOFs is drastically reduced. Since each macroelement is treated independently, the procedure is highly parallelizable and well suited for modern architectures.

In the next section, we describe how this reduced-order system integrates into the classical SIMP loop without modifying the optimization logic.

2.3 Integration Strategy within the SIMP Framework

In classical topology optimization using the SIMP method, the finite element analysis (FEA) step accounts for the majority of computational effort, especially in high-resolution or three-dimensional settings. While the optimization step itself—typically based on optimality criteria (OC) or the method of moving asymptotes (MMA)—is computationally inexpensive, the cost of solving the linear system grows rapidly with problem size. It is in this context that HiCon-FEM offers a substantial benefit by reducing the number of degrees of freedom prior to global assembly, thereby accelerating the solution of the equilibrium equations.

HiCon-FEM fits naturally into the SIMP loop. The stiffness matrix for each macroelement is projected onto a reduced basis via the operation $\tilde{\mathbf{K}}_e = \mathbf{A}_e^T\mathbf{K}_e\mathbf{A}_e$, where $\mathbf{A}_e$ encodes the hierarchical functions. This projection is followed by static condensation to eliminate internal modes, producing the final condensed matrix $\hat{\mathbf{K}}_e$. These matrices are assembled into a reduced global system, which is then solved as usual. Local displacements are subsequently recovered and used to compute compliance and sensitivities.

Given the structural inhomogeneity within macroelements, the resulting displacement fields may deviate from those predicted by classical FEM at the microscale. This discrepancy, however, is corrected by the use of filter techniques. In practice, the compatibility of HiCon-FEM with standard filter strategies—such as density and convolution filters—ensures that the design evolves smoothly and remains manufacturable. Nevertheless, in extreme cases where very few internal modes are retained and the filtering radius is too small, the optimization process may yield discontinuous or fragmented structures. Such pathological cases highlight the importance of balancing reduction aggressiveness with filtering resolution, and they reinforce the benefit of HiCon-FEM's seamless integration with well-established filtering techniques.

To reduce redundant computation during the projection step, a matrix comparison check is implemented between the current and previous iteration. For each macroelement, the difference between the new and old stiffness matrices is measured by counting the number of differing nonzero entries. If the difference falls below a fixed threshold, the projected matrix is updated incrementally via $\tilde{\mathbf{K}}_e \leftarrow \tilde{\mathbf{K}}_e^{\text{old}} + \mathbf{A}_e^T\Delta\mathbf{K}_e\mathbf{A}_e$, where $\Delta\mathbf{K}_e$ denotes the element-wise change. Otherwise, the full projection is recomputed.

This approach reduces computational overhead in later iterations when design changes become more localized.

The method has been implemented entirely in Julia, chosen for its combination of high-level syntax and low-level performance. The implementation takes advantage of parallelization, loop vectorization, sparse matrix operations, and preallocation strategies to minimize runtime and memory usage. While iterative solvers were initially considered, direct solvers proved more efficient due to the banded structure of the condensed global matrix. Despite the absence of low-level optimization or external FEM libraries, the current implementation achieves substantial performance gains, indicating that the benefit arises from the method's structure itself rather than from aggressive coding practices.

In summary, HiCon-FEM replaces only the FEA step in the SIMP loop, without altering the optimization procedure or filtering techniques. Its ability to significantly reduce the number of unknowns per iteration—combined with its parallel nature and compatibility with standard filtering—makes it an effective and scalable alternative to traditional FEM in large-scale topology optimization.

3 Numerical Results

To assess the performance and robustness of HiCon-FEM, we apply it to three classical topology optimization problems: the MBB beam, a bridge domain, and a cantilever beam. Each problem is solved using both HiCon-FEM and a classical SIMP implementation under identical physical and numerical settings, including penalization, volume fraction, and filtering strategy.

3.1 MBB Beam

The MBB beam is a classical benchmark problem in structural topology optimization, widely used for method validation due to its simplicity and well-understood response characteristics. The objective is to minimize compliance while satisfying a global volume constraint.

The design domain, boundary conditions, and external load are shown in Fig. 1. The left edge is fixed in the horizontal direction to prevent lateral motion. A vertical load is applied at the top-left corner, while a roller support is placed at the bottom-right to prevent rigid body motion. No symmetry is assumed; the full domain is optimized.

To evaluate the performance of HiCon-FEM, results are compared to those obtained using a Julia-optimized classical SIMP implementation based on the Top88 algorithm [14]. Both methods were run on the same computational setup with identical parameters: a penalization factor of $p = 3$, a volume constraint of $V/V_0 = 0.5$, and the same convergence tolerance of $\texttt{tol} = 0.16$, defined as the minimum relative change in design variables between iterations.

HiCon-FEM was tested on the MBB beam problem using two mesh resolutions:

Figure 1: Design domain, boundary conditions, and external load for the MBB beam optimization problem.

900×300 and 1200×400 elements. For both cases, a sensitivity filter was applied with radius values $r = 3$, 6, 9, and 12 to evaluate the effect of regularization on solution quality and method stability. The domain was subdivided into macroelements with varying coarseness, and identical optimization parameters were used across all runs.

In Fig. 2, we present the final optimized layouts obtained on the 1200×400 grid using a fixed macroelement division of 48×16 and varying filter radii. The visual comparison confirms that HiCon-FEM achieves results comparable to those of the classical SIMP method across a range of filtering intensities. Despite reducing internal degrees of freedom, the method preserves structural clarity and robustness throughout the convergence process.

Although the final designs produced by HiCon-FEM closely resemble those of the classical method, it is important to recall that the reduced-order solution is not the most accurate when only a small number of internal modes are retained. In this example, each macroelement uses just four internal (bubble) modes. As a result, the displacement field lacks fine-scale resolution, making the role of filtering essential to stabilize and regularize the optimization. This effect is particularly evident for the smallest filter radius ($r = 3$), where the final design becomes fragmented and ill-conditioned, with disconnected members that would be unsuitable for manufacturing. These results highlight the importance of choosing appropriate filter parameters when aggressive model reduction is used.

To isolate solver performance, Table 1 reports the runtime of the first 10 iterations for the 900×300 and 1200×400 grids under different macroelement resolutions. Each entry includes the speed-up of HiCon-FEM relative to classical SIMP and the final compliance ratio after 10 iterations.

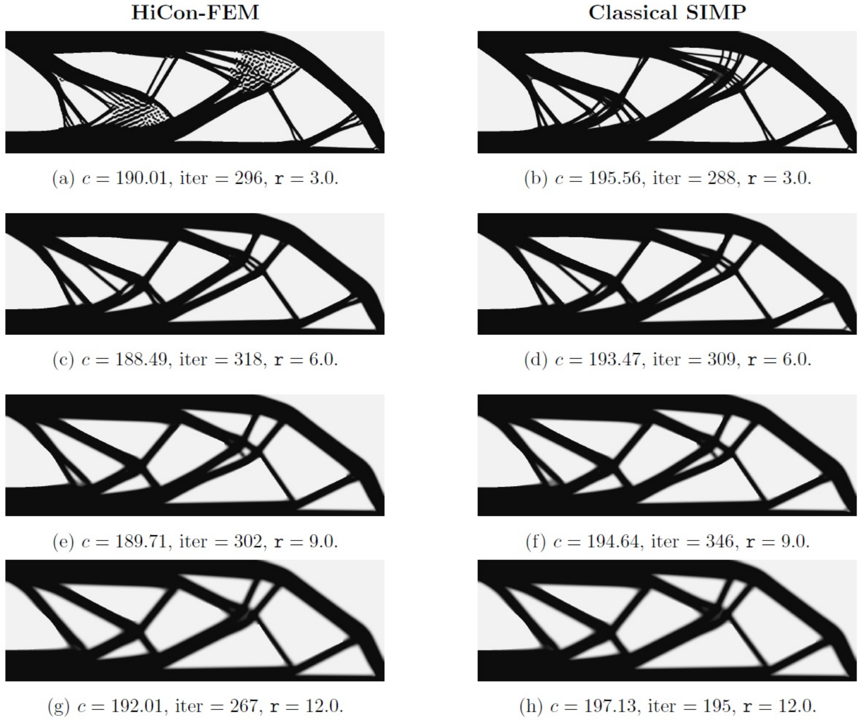

Figure 2: Final designs for the MBB beam on a 1200×400 grid using different filter radii. Left column: HiCon-FEM. Right column: Classical SIMP. Compliance (c), number of iterations, and filter radius (r) are reported for each result.

3.2 Bridge and Cantilever Problems

Given the importance of filtering in stabilizing reduced-order solutions, we conducted additional tests to verify whether the performance of HiCon-FEM holds when a different filtering technique is used. Specifically, we applied a Gaussian convolution filter, Eq. (5), to the bridge and cantilever benchmark problems. These cases allow us to evaluate whether the method remains effective under alternative smoothing strategies and design geometries.

Both the bridge and cantilever problems are solved on a 1200×400 grid with a macroelement division of 48×16. In the bridge domain, the bottom corners are fixed, and a singular vertical point load is applied at the center of the bottom edge, generating a symmetric compressive configuration that mimics typical bridge, like structural behavior. In the cantilever beam, the left edge is fully clamped, and a concentrated vertical force is applied at the midpoint of the right edge, representing a bending-dominated scenario frequently used in structural optimization benchmarks. No symmetry condition is enforced, and the full domains are included in the optimization. A Gaussian convolution filter with filtering radius of $r = 25$ is applied in both problems.

Figure 2 shows the final optimized designs for the bridge (top row) and cantilever (bottom row) problems using HiCon-FEM and classical SIMP.

Resolution	Classical Time (10 iter)	Macroelements	HiCon Time (10 iter)	Speed-Up
900×300	59.52 s	30×10	26.61 s	2.24×
		45×15	27.19 s	2.19×
		60×20	32.84 s	1.81×
		75×25	51.70 s	1.15×
1200×400	102.94 s	30×10	55.52 s	1.85×
		48×16	51.53 s	2.00×
		50×20	51.81 s	1.99×
		60×20	55.33 s	1.86×

Table 1: Runtime for 10 iterations of HiCon-FEM vs. classical SIMP on the MBB beam problem under different macroelement resolutions.

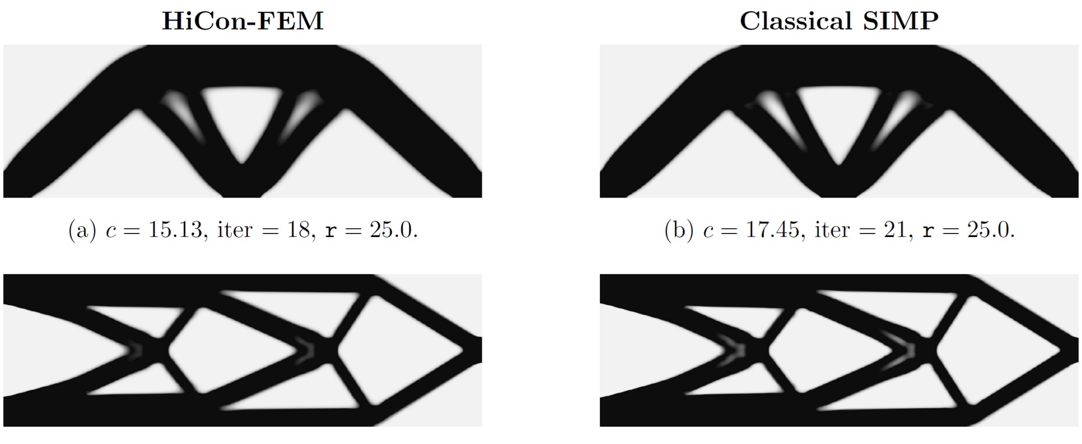

(a) $c = 15.13$, iter $= 18$, $\mathtt{r} = 25.0$.

(b) $c = 17.45$, iter $= 21$, $\mathtt{r} = 25.0$.

(c) $c = 172.82$, iter $= 53$, $\mathtt{r} = 25.0$.

(d) $c = 173.01$, iter $= 66$, $\mathtt{r} = 25.0$.

Figure 3: Comparison of final optimized designs using HiCon-FEM (left column) and classical SIMP (right column) with a Gaussian convolution filter ($r = 25.0$). Top row: bridge problem. Bottom row: cantilever problem. Compliance (c), number of iterations, and filter radius are annotated. The stopping criterion was set to `tol` = 0.16.

In both cases, the results are remarkably similar in topology and compliance, demonstrating that HiCon-FEM remains accurate even under a different filtering strategy. While HiCon-FEM slightly underestimates the compliance compared to the classical method, it requires fewer iterations to converge, highlighting its potential for accelerating large-scale optimizations without sacrificing structural fidelity.

4 Conclusions

This work introduced a hierarchical condensation-based finite element stage designed to accelerate topology optimization. The method integrates seamlessly into the topology optimization framework, requiring no modification to the other optimization steps and operating as a drop-in replacement for standard FEA.

Through numerical experiments on MBB beam, bridge, and cantilever problems, we demonstrated that the method significantly reduces computation time while maintaining high-quality results. For moderate macroelement sizes and appropriate filter parameters, speed-ups of up to 2.24$\times$ were observed with minimal deviation in compliance.

While the reduced-order model offers substantial computational benefits, its accuracy depends on both the number of internal modes retained and the strength of the applied filter. In particular, the use of filtering proved essential to compensate for the localized approximation introduced by hierarchical projection.

Future developments will focus on extending the capabilities to handle three-dimensional domains and nonlinear material behavior, where gains are expected to be more important for these cases in which reducing the global system size is critical.

Acknowledgements

VY, FJM and NHK acknowledge the funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 956401.

References

[1] M.P. Bendsøe, N. Kikuchi, "Generating optimal topologies in structural design using a homogenization method", Computer Methods in Applied Mechanics and Engineering, 71(2), 197–224, 1988.

[2] M.P. Bendsøe, "Optimal shape design as a material distribution problem", Structural Optimization, 1(4), 193–202, 1989.

[3] M.P. Bendsøe, O. Sigmund, Topology Optimization: Theory, Methods, and Applications, Springer, 2004.

[4] J.-H. Zhu, W.-H. Zhang, D. Kang, "Topology Optimization in Aircraft and Aerospace Structures Design", Archives of Computational Methods in Engineering, 23(4), 595–622, 2016.
[5] R.J. Yang, A.I. Chahande, "Automotive applications of topology optimization", Structural Optimization, 9(3), 245–249, 1995.
[6] L.L. Beghini, A.M. Carrion, S. Guest, "Connecting architecture and engineering through structural topology optimization", Engineering Structures, 59, 716–726, 2014.
[7] J. Wu, L. Wang, X. Zhang, "A System for High-Resolution Topology Optimization", IEEE Transactions on Visualization and Computer Graphics, 22(3), 1195–1208, 2016.
[8] L. Xia, P. Breitkopf, "Concurrent topology optimization design within FE^2 multiscale analysis framework", Computer Methods in Applied Mechanics and Engineering, 278, 524–542, 2014.
[9] Z. Wu, Y. Wang, K. Wang, "Topology optimization of hierarchical lattice structures with substructuring", Computer Methods in Applied Mechanics and Engineering, 345, 602–617, 2019.
[10] D.A. White, T. Zhang, J. Arróyave, "Multiscale topology optimization using neural network surrogate models", Computer Methods in Applied Mechanics and Engineering, 346, 1118–1135, 2019.
[11] S. Park, M. Park, S. Lee, "Structural Topology Optimization Using Two-Level Dynamic Condensation Scheme", Proceedings of the 47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2006.
[12] H.-G. Kim, J. Kim, C.-S. Han, "Structural topology optimization based on system condensation", Finite Elements in Analysis and Design, 92, 26–35, 2014.
[13] H.S. Koh, Y.S. Kim, J.Y. Lee, "Efficient topology optimization using substructuring-based model order reduction method", Computers & Structures, 228, 106146, 2020.
[14] E. Andreassen, A. Clausen, M. Schevenels, B.S. Lazarov, O. Sigmund, "Efficient topology optimization in MATLAB using 88 lines of code", Struct. Multidisc. Optim., 43, 1–16, 2011.