

HIERARCHICAL HIGH-ORDER BEAM MODELS FOR OPTIMISATION-ORIENTED ANALYSIS OF COMPOSITE WING-LIKE STRUCTURES

Donato Cardone[1], Rauno Cavallaro[2] & Andrea Cini[1]

[1]Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, 28911, Spain
[2]Barcelona Supercomputing Center (BSC), Barcelona, 08034, Spain

Abstract

This paper investigates hierarchical high-order beam models based on the Carrera Unified Formulation for the analysis and size optimisation of laminated composite wing-like structures. The objective is to identify a compact one-dimensional finite-element model able to capture the relevant sectional response at a computational cost compatible with repeated preliminary-design evaluations. Two benchmarks are considered: a laminated flat panel subjected to axial, bending, and torsional loads, and a non-stiffened composite box beam subjected to an eccentric tip load producing coupled bending and torsion. Taylor- and Lagrange-expansion CUF models are compared with shell and solid finite-element references in terms of strain energy, displacement response, local stresses, degrees of freedom, runtime, and optimisation performance. For the flat-panel benchmark, the fourth-order Taylor-expansion model (TE4) keeps the strain-energy error below 0.5% for all load cases while reducing runtime by more than 99.8%. For the box-beam benchmark, TE8 gives the most accurate global response, whereas TE4 provides the best accuracy-cost compromise, with global errors close to 2.6%, 1395 degrees of freedom, and a runtime of approximately 0.11 s. The selected TE4 model is then used in a preliminary sizing study, achieving a 46.55% normalised mass reduction in 10 iterations. Results show that moderate-order CUF beam models are effective for preliminary composite wing-box sizing, while higher-fidelity models remain necessary for detailed local-stress verification.

Keywords: Carrera Unified Formulation; high-order beam models; composite structures; wing-box; sectional mechanics

1. Introduction

Composite materials are widely used in modern aircraft structures as they provide high specific stiffness and strength, together with the possibility of tailoring their mechanical response through ply orientation and stacking sequence. In wing structures, this tailoring capability is especially relevant because skins, spars, ribs, and stiffened panels must simultaneously satisfy different stiffness and structural integrity requirements ensuring also mass efficiency [1, 2, 3]. However, the same anisotropic and heterogeneous nature making laminated composites attractive also increases the modelling complexity, particularly when local through-thickness stresses, shear effects, warping, and bending-torsion coupling are relevant.

High-fidelity finite-element models, based on three-dimensional solid discretisations can predict detailed stress and strain fields, but their computational cost becomes prohibitive when many repeated analyses are required. This is the case in preliminary sizing, design-space exploration, gradient-based optimisation, and model-order selection studies. Shell models are therefore commonly adopted in industrial structural simulations as they offer a practical compromise between modelling cost and structural fidelity. Nevertheless, for thin-walled laminated components and closed-section composite

structures, shell models may still miss relevant sectional effects associated with through-thickness response, transverse shear and local skin-web stress redistribution [4, 5, 6, 7].

Aircraft wing-box optimisation has traditionally relied on reduced structural modelling, shell-based GFEM, equivalent-plate formulations, or panel-level sizing methods to keep the design problem computationally tractable [8, 9, 10]. These approaches have demonstrated that substantial mass reduction can be achieved when structural sizing is formulated as an optimisation problem. However, for composite wing structures, the choice of the structural model remains critical: an extremely simple model may lead to inaccurate strain or stress constraints, whereas a very detailed model may become impractical for repeated gradient-based design evaluations [11]. This motivates the development of intermediate-order models characterising the essential mechanical response of composite structures while remaining computationally inexpensive.

The Carrera Unified Formulation (CUF) provides a systematic framework to derive hierarchical one-dimensional, two-dimensional, and three-dimensional structural theories with variable kinematic richness [12, 13, 14, 15]. In the beam context, the three-dimensional displacement field is expanded over the cross-section, allowing classical and refined beam theories to be obtained within the same finite-element framework. Taylor-expansion models provide compact equivalent-single-layer descriptions, while Lagrange-expansion models can introduce more local and layer-wise sectional refinement [12, 16, 17]. This hierarchy makes CUF particularly attractive for optimisation analysis, since the model order can be selected according to the required trade-off between global accuracy, local sectional fidelity, and computational cost.

Recent works have proved the potential of CUF-based models for aerospace structural analysis and optimisation. In particular, CUF beam models have been coupled with gradient-based optimisation strategies for composite aircraft components and wing-like structures [18, 19]. Those studies demonstrated the feasibility of using refined beam models inside optimisation loops, but they also raised an open question: which level of the CUF hierarchy is sufficiently detailed to capture the relevant sectional mechanics of composite wing structures while remaining computationally efficient enough for repeated sizing analyses? Addressing this question requires not only a comparison of global displacements or strain energy, but also a local assessment of strain in composite sections.

The present work addresses this model-order selection gap for optimisation-oriented composite wing-box analysis. Its novel aspect is the quantitative identification of a production-order CUF beam model by combining, within the same study, global stiffness accuracy, local stress-recovery capability, computational cost, and preliminary sizing performance. Two numerical benchmarks of increasing structural complexity are considered. The first benchmark is a laminated flat panel subjected to elementary axial, bending, and torsional load resultants. This test case isolates the main deformation mechanisms and clarifies how increasing CUF order affects global stiffness and through-thickness stress recovery. The second benchmark is an unstiffened closed-section composite box beam subjected to an eccentric vertical tip load, producing coupled bending and torsion representative of simplified wing-box behaviour. CUF Taylor and Lagrange models are compared with shell and solid finite-element references in terms of strain energy, dominant displacement components, local stress recovery, degrees of freedom, runtime, and optimisation performance.

The main contributions of the paper are: (i) a mechanism-by-mechanism assessment of the CUF hierarchy for axial, bending, and torsional response of a laminated panel; (ii) a closed-section box-beam benchmark testing whether the same model-order trends persist under coupled bending-torsion; (iii) a quantitative accuracy–cost comparison against shell and solid finite-element references; and (iv) a preliminary thickness-sizing study using the selected production-order CUF model. The objective is not to replace high-fidelity verification models, but to identify a compact beam model suitable for early-stage and optimisation-oriented composite sizing. The results show that TE4 provides a practical accuracy–cost compromise for repeated design evaluations, while higher-order CUF and three-dimensional models remain necessary for detailed local-stress verification and final design assessment.

2. High-Order CUF Beam Modelling Framework

CUF derives refined beam theories by expanding the three-dimensional displacement field over the cross-section rather than prescribing a fixed beam kinematics [12, 13, 14]. The displacement approx-

imation is written as

$$\mathbf{u}(x,y,z) = \sum_{\tau=1}^{M} F_\tau(y,z)\,\mathbf{u}_\tau(x), \tag{1}$$

where F_τ are sectional expansion shape functions, $\mathbf{u}_\tau$ are the generalized displacement variables along the beam axis, and M controls the richness of the sectional kinematics. Increasing the expansion order allows the model to recover shear deformation, warping, sectional distortion, bending-torsion coupling, and non-uniform strain distributions.

2.1 Taylor and Lagrange expansion kinematics

Taylor-expansion (TE) models use polynomial functions defined over the full cross-section and provide a compact equivalent-single-layer description [12, 13, 14]. For a Taylor expansion of order N, the number of sectional terms is

$$M = \frac{(N+1)(N+2)}{2}. \tag{2}$$

Since three displacement components are associated with each sectional term, the number of generalized unknowns per beam node is $3M$. Increasing N therefore enriches the global sectional kinematics while preserving a one-dimensional finite-element discretisation along the beam axis. For this reason, TE models are the primary candidates for optimisation-oriented analyses in this work.

Lagrange-expansion (LE) models adopt a local sectional refinement strategy. The cross-section is discretised into two-dimensional Lagrange sectional elements, here based on nine-node L9 elements. Each sectional node carries the three displacement components of the three-dimensional field, allowing the deformation to vary locally across material interfaces, wall boundaries, corners, and geometric discontinuities. LE models are therefore retained as refined CUF references for local and layer-wise stress recovery, although their higher number of sectional unknowns makes them less suitable for repeated optimisation analyses.

Figure 1 summarises the two sectional-enrichment strategies adopted in this work. TE models increase the global polynomial order over the complete section, whereas LE models increase the local cross-sectional resolution through L9 sectional elements.

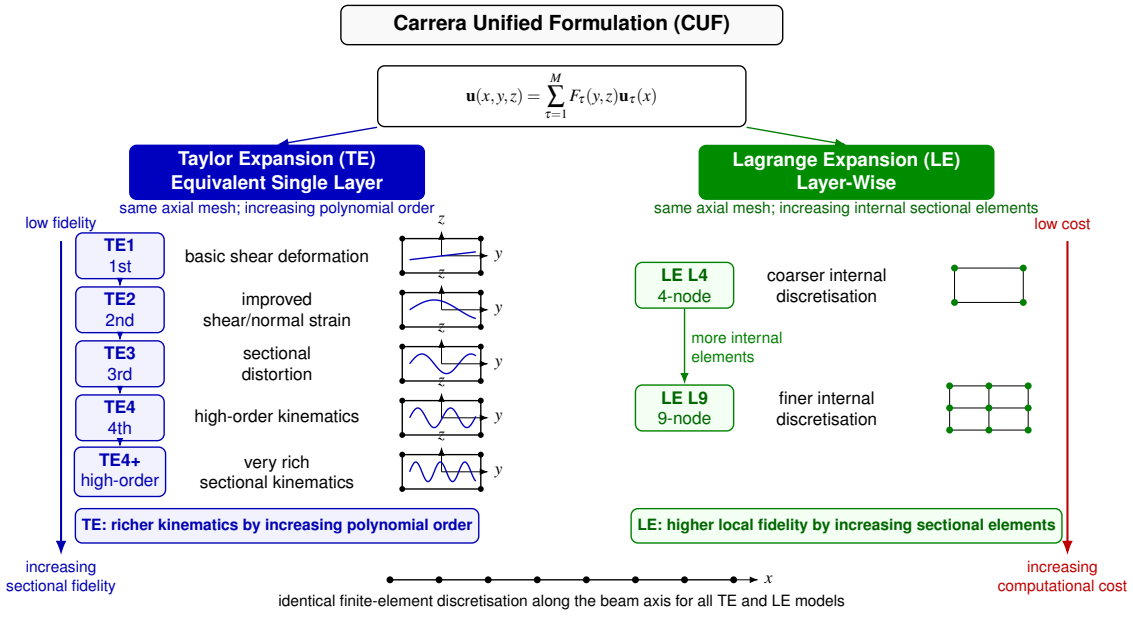


Figure 1 – CUF beam-model hierarchy adopted in this work: Taylor-expansion models provide compact global sectional enrichment, whereas Lagrange-expansion models provide local sectional refinement.

2.2 Finite Element Approximation and Laminate Representation

The generalized displacement variables are discretized along the beam axis with third-order, four-node beam elements for all CUF models, for the sake of isolating the influence of sectional kinematic enrichment during comparisons. The analysed laminates are modelled stacking orthotropic unidirectional carbon fibre reinforced epoxy plies with prescribed material orientation. TE models are used as cost-effective reduced order models, whereas LE models are retained as refined CUF references.

3. Numerical Test Cases

3.1 Material Properties and Laminate Definition

All tested configurations employ the same carbon-fibre-reinforced laminate to isolate the effect of structural configuration and model order. Table 1 lists the in-plane mechanical properties of a lamina.

Table 1 – Orthotropic material properties adopted in the numerical analyses.

Property	Value	Unit
E_1	130	GPa
E_2	10	GPa
E_3	10	GPa
ν_{12}	0.25	-
ν_{13}	0.25	-
ν_{23}	0.35	-
G_{12}	5	GPa
G_{13}	5	GPa
G_{23}	3.7	GPa
ρ	1600	kg/m^3

The laminate is symmetric and balanced, showing the following stacking sequence

$$[0/90/+45/-45]_s, \tag{3}$$

with 0.25 mm thick plies for a total laminate thickness of 2.0 mm. The symmetric balanced stacking sequence avoids bending–extension and shear–extension coupling.

3.2 TC1: Laminated Flat Panel Under Elementary Loads

The Test Case 1, TC1, is a clamped flat laminated panel used to characterise elementary deformation mechanisms. Table 2 lists the panel dimensions. Axial loading, pure bending, and pure torsion are applied at the free-end section using distributed load resultants to reduce local load-introduction effects.

Table 2 – Geometrical properties of the TC1 flat-panel benchmark.

Quantity	Value	Unit
Length L	1000	mm
Width b	250	mm
Thickness t	2.0	mm
Ply thickness	0.25	mm
Laminate	$[0/90/+45/-45]_s$	-

3.2.1 TC1 load cases and finite-element modelling details.

Figure 2 shows the TC1 benchmark configuration and the three elementary load cases (LC) considered for the laminated panel. The panel is fully clamped at the root section, while the opposite edge is left free and loaded through equivalent section resultants. This boundary-condition definition is used

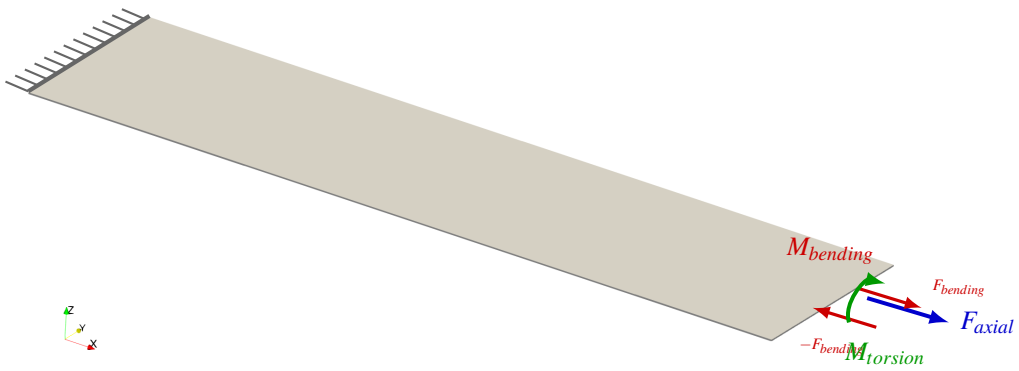


Figure 2 – TC1 flat-panel benchmark with clamped root and axial, bending, and torsional load resultants applied at the free-end section.

consistently for the CUF beam models and for the shell and solid finite-element references, so that differences in the response can be attributed to the modelling assumptions rather than to the load or constraint implementation.

The three load cases are selected to activate progressively different structural mechanisms. The axial case (LC1) applies a resultant force $F_x = 10{,}000$ N at the free-end section and mainly assesses membrane stiffness and the recovery of the axial stress component σ_{xx}. The pure-bending case (LC2) applies a moment $M_y = 20$ N m, generated by an axial force couple at the free-end section, and introduces a through-thickness axial-stress gradient. The torsional case (LC3) applies a resultant moment $M_x = 20$ N m about the beam axis and activates shear-dominated stress states, sectional warping, and non-uniform through-thickness stress distributions. This last case is therefore the most demanding TC1 configuration for reduced-order beam models.

The shell and solid reference models use the mesh hierarchy reported in Table 4. The finest solid model, H8-4, is adopted as the TC1 reference solution for global and local error evaluations.

3.3 TC2: Composite Box Beam Under Coupled Bending and Torsion

TC2 extends the analysis to a closed thin-walled composite box beam representative of a simplified wing-box. Differently from TC1, which isolates elementary deformation mechanisms, TC2 combines bending, torsion, skin-web load transfer, restrained warping, and sectional stiffness redistribution in a single structural benchmark. Figure 3 shows the geometry, boundary conditions, and loading configuration adopted for this test case.

The box beam has a rectangular closed section with no stiffeners. Its dimensions are reported in Table 3: the beam is $L = 1200$ mm long, $b = 120$ mm wide, and $h = 60$ mm high, with a uniform wall thickness of $t = 2.0$ mm. All four walls are modelled using the same symmetric balanced laminate adopted in TC1, namely $[0/90/+45/-45]_s$. The root section at $x = 0$ is fully clamped, while the section at $x = L$ is free. This boundary-condition setup reproduces a cantilevered wing-box-like configuration while preserving a slender beam behaviour.

Table 3 – Geometrical properties of the TC2 composite box-beam benchmark.

Quantity	Value	Unit
Length L	1200	mm
Width b	120	mm
Height h	60	mm
Wall thickness t	2.0	mm
Laminate	$[0/90/+45/-45]_s$	-

As shown in Figure 3, an eccentric vertical force $F_z = 1000$ N is applied at the free-end section, at the lateral edge of the box section. The load is applied at $y = +b/2$ and $z = 0$, generating both a bending moment and a torsional moment. The corresponding resultants are

$$M_y = F_z L = 1200 \text{ N m}, \qquad M_x = F_z \frac{b}{2} = 60 \text{ N m}. \tag{4}$$

This loading condition is therefore more representative of a simplified wing-box response than a centred vertical force, because it activates the coupled bending-torsion behaviour of the closed section. The shell and solid reference models follow the refinement hierarchy reported in Table 4. For TC2, the finest available solid model, H8-3, is adopted as the reference solution for the global response assessment.

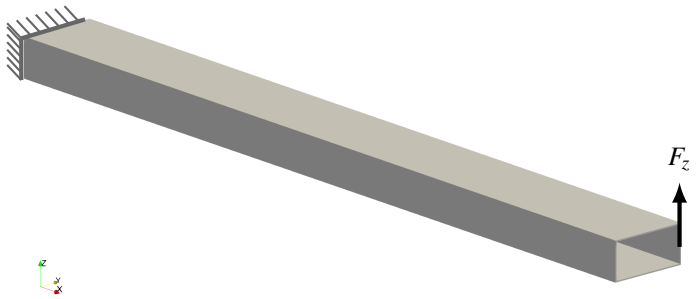


Figure 3 – TC2 composite box-beam benchmark configuration with fully clamped root section and eccentric tip loading.

3.4 Comparison and validation of numerical models

The numerical comparison assesses both accuracy and computational efficiency. CUF beam models with different sectional kinematics are compared against independent shell and solid finite-element models. Shell models are used as an industrial baseline, whereas solid models provide the three-dimensional references for global quantities and representative local stress distributions.

The reference shell and solid mesh hierarchy is summarised in Table 4. CUF models are not included in the table because they use the same one-dimensional mesh along the span and differ only in their sectional description. In TC1, Taylor-expansion models up to TE10 are analysed together with selected LE-L9 models. In TC2, Taylor-expansion models up to TE8 are compared with the shell and solid refinement levels. The summary tables retain only representative models: TE4 as the production-order candidate, TE8 as a higher-order Taylor model, selected LE models as refined CUF references, the finest shell models as industrial baselines, and the finest solid models as reference solutions.

Global accuracy is evaluated primarily through the total strain energy, while dominant displacement components are retained as secondary engineering indicators. A Cartesian reference frame is adopted in both test cases, with the x-axis aligned with the panel or beam span and $x = 0$ located at the clamped root section. For TC1, y is aligned with the panel width and z is the through-thickness coordinate. For TC2, y and z are aligned with the box width and height, respectively, with the origin at the centroid of the closed section.

Local quantities are extracted away from boundary constraints and load-introduction regions. For TC1, through-thickness stress profiles are evaluated at $x = 0.5L$ and $y = 0$. For TC2, local sectional stresses are evaluated at $x = 0.5L$, with the corner comparison performed in the lower-left skin–web junction. The relative error of a response quantity q is calculated as

$$e_q = \frac{|q_{\text{model}} - q_{\text{ref}}|}{|q_{\text{ref}}|} \times 100. \tag{5}$$

Runtime and number of degrees of freedom are used to quantify computational efficiency.

Table 4 – Reference finite-element mesh hierarchy for TC1 and TC2.

Test case	Model	Mesh level	Δx [mm]	Δy [mm]	Laminate/thickness model	Approx. elements
TC1	H8 solid	TC1-H8-1 coarse	25.0	12.5	1 element/ply = 8	6,400
TC1	H8 solid	TC1-H8-2 medium	12.5	6.25	1 element/ply = 8	25,600
TC1	H8 solid	TC1-H8-3 fine	6.25	3.125	1 element/ply = 8	102,400
TC1	H8 solid	TC1-H8-4 fine	5.0	2.5	1 element/ply = 8	160,000
TC1	PCOMPP shell	TC1-SH-1 coarse	25.0	12.5	Ply-based laminate	800
TC1	PCOMPP shell	TC1-SH-2 medium	12.5	6.25	Ply-based laminate	3,200
TC1	PCOMPP shell	TC1-SH-3 fine	6.25	3.125	Ply-based laminate	12,800
TC1	PCOMPP shell	TC1-SH-4 fine	5.0	2.5	Ply-based laminate	20,000
TC2	H8 solid	TC2-H8-1 coarse	40	10	1 element/ply = 8	8,640
TC2	H8 solid	TC2-H8-2 medium	20	5	1 element/ply = 8	34,560
TC2	H8 solid	TC2-H8-3 fine	10	2.5	1 element/ply = 8	138,240
TC2	PCOMPP shell	TC2-SH-1 coarse	40	10	Ply-based laminate	1,080
TC2	PCOMPP shell	TC2-SH-2 medium	20	5	Ply-based laminate	4,320
TC2	PCOMPP shell	TC2-SH-3 fine	10	2.5	Ply-based laminate	17,280
TC2	PCOMPP shell	TC2-SH-4 fine	5	1.25	Ply-based laminate	69,120

3.5 Computational setup and finite-element implementation

The shell and solid reference models were generated in HyperMesh 2024 and solved with OptiStruct 2024 using linear static analyses. CUF analyses were performed with an in-house C++ code executed under Linux through WSL. All runtime comparisons were obtained on the same workstation, equipped with an Intel Core i9-14900HX CPU and 32 GB RAM, using CPU-based computations only.

The same orthotropic ply properties, stacking sequence, thickness definitions, boundary conditions, and equivalent end-section resultants were used for the CUF, shell, and solid models. Shell references use structured quadrilateral meshes with ply-based laminate properties, whereas solid references use structured eight-node hexahedral meshes with one element per ply through the laminate thickness. In the CUF models, all Taylor and Lagrange expansion orders use the same spanwise discretisation, consisting of ten third-order four-node beam elements unless otherwise stated. Therefore, differences among CUF models are mainly due to sectional kinematic enrichment rather than spanwise mesh refinement.

3.6 Mesh-sensitivity assessment and reference-model selection

Before comparing the CUF beam models with the shell and solid solutions, a mesh-sensitivity assessment is performed for the reference finite-element models. The objective of this step is not to prove full pointwise stress convergence, which is particularly demanding in laminated composites, but to identify sufficiently refined shell and solid models for a consistent evaluation of global response quantities and representative local stress distributions.

For both TC1 and TC2, the shell and solid models are progressively refined in the in-plane directions, as reported in Table 4. In the solid models, one hexahedral element is retained through each ply thickness, so that the through-thickness laminate discretisation is kept physically consistent while the in-plane mesh size is reduced. The total strain energy and the dominant displacement component are used as the primary mesh-sensitivity indicators because they are integral or global quantities and are therefore less affected by local numerical oscillations than pointwise stresses.

The finest solid models, H8-4 for TC1 and H8-3 for TC2, are adopted as the reference solutions for the error evaluations. These models are selected because they represent the most refined available three-dimensional descriptions of the corresponding benchmarks. However, local stress quantities remain more mesh-sensitive than strain energy and displacements, especially near constrained boundaries, load-application regions, ply interfaces, and skin-web junctions. For this reason, local stress comparisons are extracted away from boundary and load-introduction regions: at $x = 0.5L$ and $y = 0$ for TC1, and at the mid-span section for TC2. The local stress results are therefore interpreted as representative sectional-recovery indicators rather than as mesh-independent pointwise peak-stress predictions.

4. TC1 Results: Laminated Panel Mechanical Response

The TC1 results are used to assess how the CUF hierarchy behaves under elementary deformation mechanisms. Each load case is evaluated through the strain-energy error with respect to the H8-4

solid reference and through a representative through-thickness stress distribution at $x = 0.5L$. This combined assessment avoids selecting a reduced-order model only from global stiffness quantities, since similar strain-energy errors may still correspond to different local stress fields.

4.1 LC1: Axial Loading

The axial load case is the least demanding TC1 configuration because the response is mainly governed by the laminate membrane stiffness. As shown in Figure 4, the TE sequence reaches very small strain-energy errors already at low order, and further sectional enrichment provides only marginal improvement at increased computational cost. The LE models are also accurate, but their local sectional refinement does not provide a significant advantage for this membrane-dominated response.

The local σ_{xx} distribution in Figure 5 confirms that the selected CUF models reproduce the ply-wise axial-stress pattern of the H8-4 reference along the thickness path. Therefore, LC1 establishes the lower-complexity limit of the hierarchy: low- and moderate-order TE models are already sufficient for axial stiffness and axial-stress recovery.

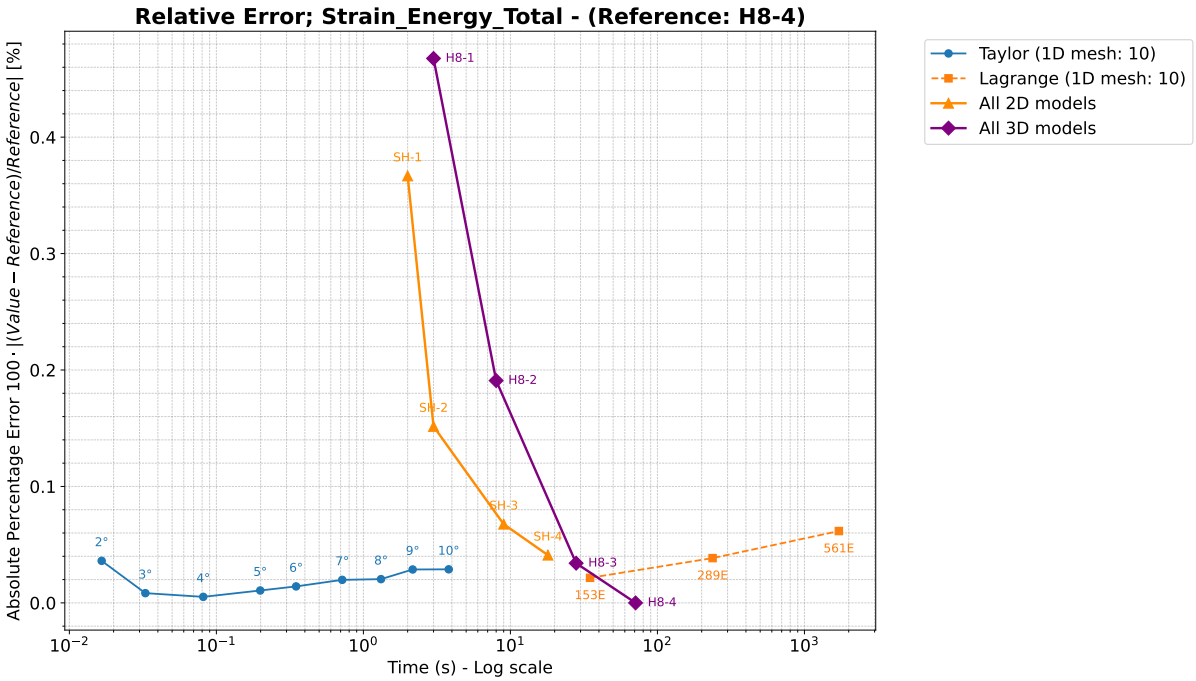
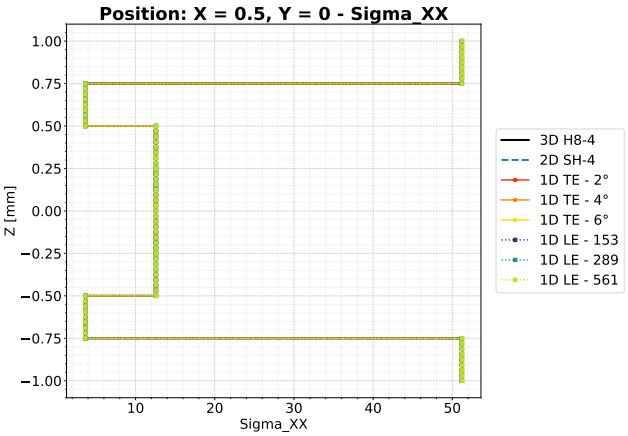


Figure 4 – Relative error in total strain energy for TC1-LC1 axial loading as a function of computational time. The H8-4 solid model is used as reference. The labels reported in the graph identify the CUF model order: TE labels indicate the Taylor expansion order, whereas LE labels indicate Lagrange-expansion models based on local sectional discretisation.

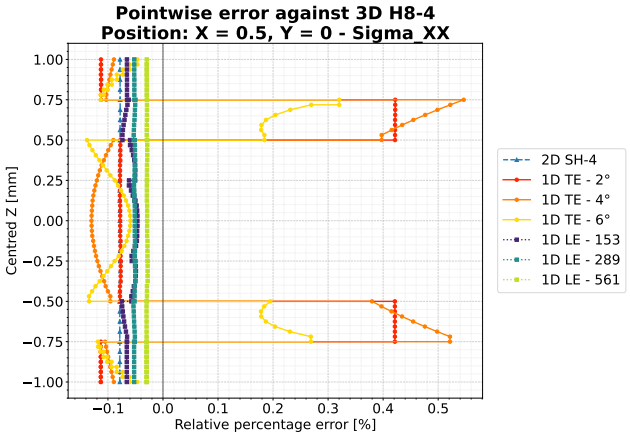
4.2 LC2: Pure Bending

Pure bending is more discriminating than axial loading because it generates a non-uniform σ_{xx} distribution through the laminate thickness. Figure 6 shows that TE4, TE8, the selected LE model, and SH-4 all provide comparable global stiffness accuracy. In particular, TE4 gives $e_U = 0.456\%$, while the selected LE model gives $e_U = 0.274\%$.

The through-thickness comparison in Figure 7 shows that the selected CUF models reproduce the bending-driven sign change and the main ply-wise stress discontinuities. The largest discrepancies remain concentrated near ply interfaces and in regions with steeper stress gradients. Thus, TE models provide a favourable accuracy–cost compromise for bending, whereas LE models are mainly useful when more detailed local stress recovery is required.



(a) Through-thickness axial stress distribution.



(b) Pointwise relative error with respect to H8-4.

Figure 5 – Through-thickness recovery of σ_{xx} for TC1-LC1 axial loading at $x = 0.5L$ and $y = 0$. The H8-4 solid model is adopted as the reference solution for the pointwise error evaluation.

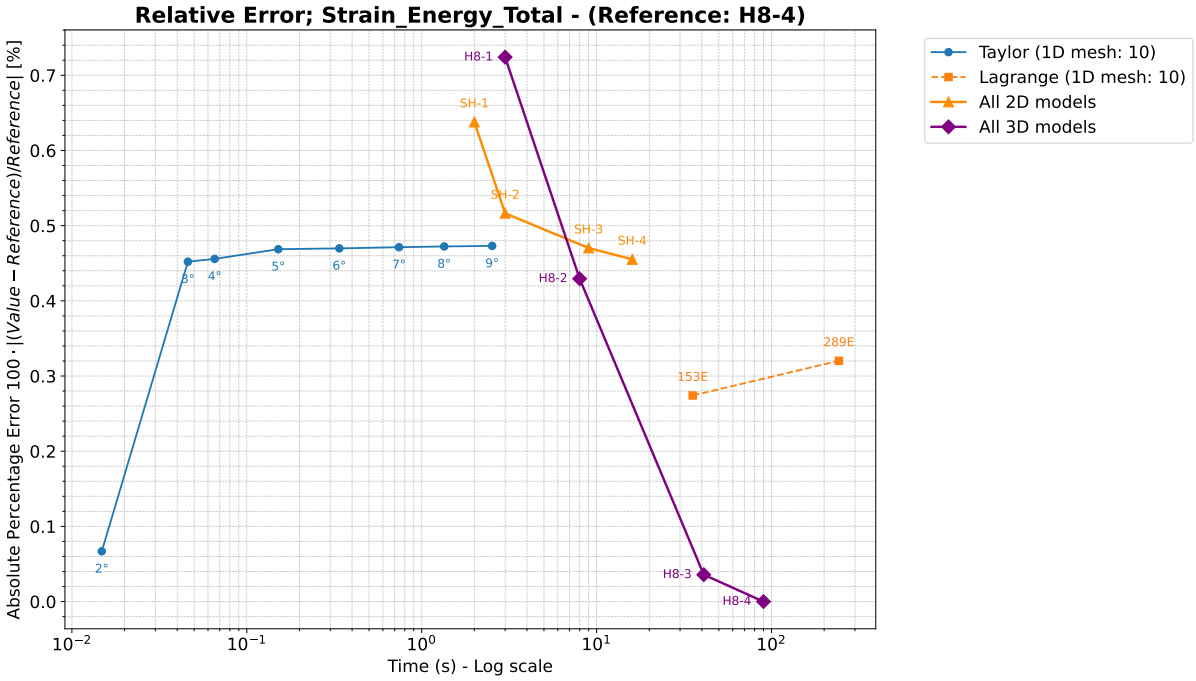
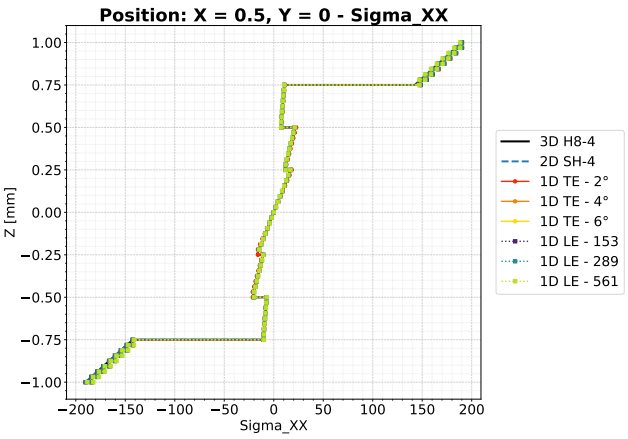
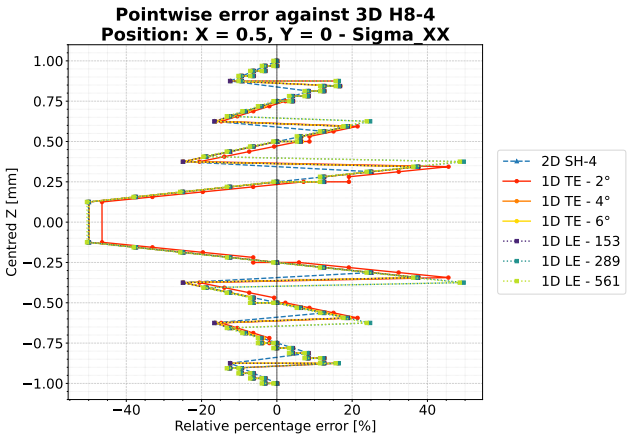


Figure 6 – Relative error in total strain energy for TC1-LC2 pure bending as a function of computational time. The H8-4 solid model is used as reference. The same plotting convention adopted for LC1 is used: TE labels indicate the Taylor expansion order, LE labels identify Lagrange-expansion models, and shell/solid models are included as finite-element comparison levels.



(a) Through-thickness axial stress distribution.



(b) Filtered pointwise relative error with respect to H8-4.

Figure 7 – Through-thickness recovery of σ_{xx} for TC1-LC2 pure bending at $x = 0.5L$ and $y = 0$. The H8-4 solid model is adopted as the reference solution for the pointwise error evaluation.

4.3 LC3: Pure Torsion

Torsion is the most severe TC1 case because it activates shear-dominated response and sectional warping. Figure 8 shows that the strain-energy error decreases more gradually with Taylor order than in LC1 and LC2.

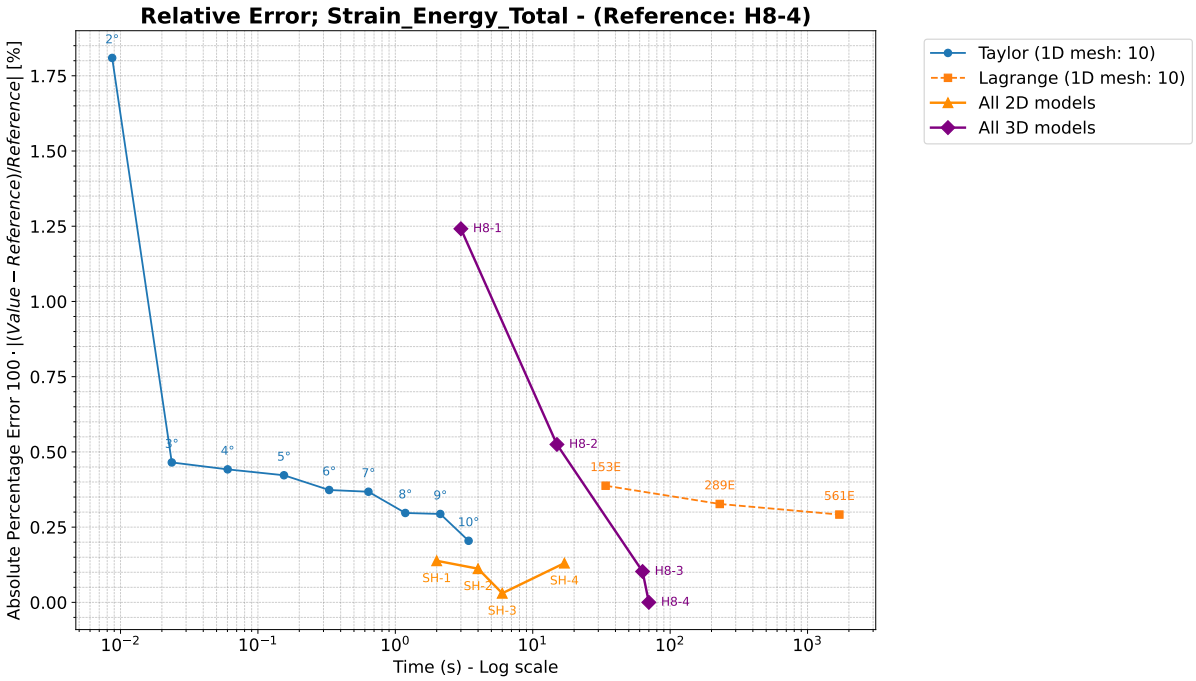
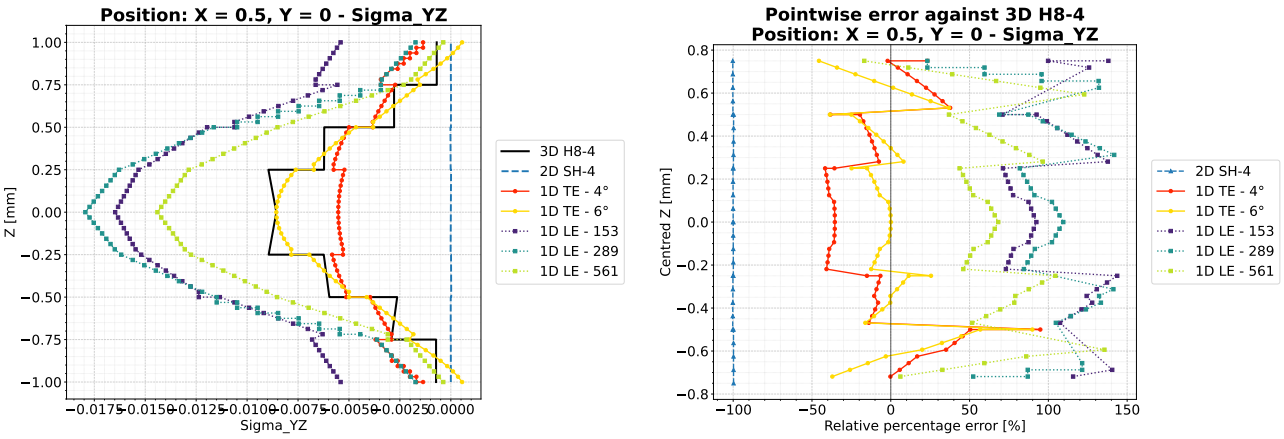


Figure 8 – Total strain-energy error for TC1-LC3 pure torsion as a function of computational time. The H8-4 solid model is adopted as the reference solution.

The clearest local indicator is σ_{yz}, shown in Figure 9. The SH-4 shell model remains efficient for global quantities but does not reproduce the through-thickness shear-stress profile of the H8-4 reference. TE4 already captures the main non-zero shear-stress trend and substantially improves the local stress prediction with respect to the shell model, although higher-order TE and LE models may still be required for more accurate local-stress verification.



(a) Through-thickness distribution of σ_{yz}.

(b) Pointwise relative error with respect to H8-4.

Figure 9 – Sectional recovery of σ_{yz} for TC1-LC3 pure torsion at $x = 0.5L$ and $y = 0$. The comparison highlights the capability of the TE4 model to recover a three-dimensional-like through-thickness shear-stress distribution more accurately than the SH-4 shell model.

Table 5 – Summary of TC1 global response, computational cost, and errors with respect to the selected H8-4 solid reference.

Load case	Model	DOFs	Runtime [s]	Time saving [%]	U	e_U [%]	$u_{\mathrm{dom,max}}$	$e_{u_{\mathrm{dom,max}}}$ [%]
LC1 axial	H8-4 ref.	545406	71.000	-	1964.498	-	0.399	-
LC1 axial	TE4	1395	0.081	99.89	1964.599	0.005	0.393	1.444
LC1 axial	TE8	4185	1.318	98.14	1964.898	0.020	0.393	1.426
LC1 axial	LE-L9	14229	34.814	50.97	1964.921	0.022	0.393	1.422
LC1 axial	SH-4	243006	18.000	74.65	1965.303	0.041	0.398	0.228
LC2 bending	H8-4 ref.	545400	90.000	-	14508.355	-	746.560	-
LC2 bending	TE4	1395	0.066	99.93	14574.482	0.456	749.975	0.457
LC2 bending	TE8	4185	1.345	98.51	14576.887	0.472	750.035	0.465
LC2 bending	LE selected	14229	35.425	60.64	14548.124	0.274	749.034	0.331
LC2 bending	SH-4	243006	16.000	82.22	14574.389	0.455	750.700	0.555
LC3 torsion	H8-4 ref.	550854	70.000	-	31885.082	-	485.796	-
LC3 torsion	TE4	1395	0.060	99.91	31744.167	0.442	484.575	0.251
LC3 torsion	TE8	4185	1.182	98.31	31790.372	0.297	485.090	0.145
LC3 torsion	LE selected	14229	33.912	51.55	31761.484	0.388	484.598	0.247
LC3 torsion	SH-4	243006	17.000	75.71	31926.636	0.130	486.796	0.206

4.4 Synthesis of the TC1 Model results

Table 5 summarises the TC1 comparison using representative models. TE4 is retained as the computationally efficient production-order candidate, TE8 quantifies the effect of additional Taylor enrichment, the selected LE model represents a refined CUF alternative, and SH-4 provides the converged shell baseline. All errors are computed with respect to the corresponding H8-4 solid reference.
The time saving is computed as

$$\Delta t_{\mathrm{save}} = 100\left(1 - \frac{t_{\mathrm{model}}}{t_{\mathrm{ref}}}\right), \tag{6}$$

where t_{model} and t_{ref} are the measured wall-clock solution times of the analysed model and of the corresponding H8 solid reference, respectively. The reported time includes the linear static solution phase and the recovery of the quantities used in the comparison, but excludes pre-processing, mesh generation, and post-processing operations.
Across the three load cases, TE4 keeps the strain-energy error below 0.5% while reducing the measured solution time by more than 99.8% with respect to the H8 solid reference. TE8 improves the torsional response but at higher computational cost, whereas LE-L9 provides refined sectional recovery with runtimes closer to verification models than to production-order optimisation models.
The main conclusion from TC1 is therefore mechanism-dependent. Axial loading requires little sec-

tional enrichment, bending benefits from moderate enrichment, and torsion is the discriminating case for model-order selection. TE4 is retained as the production-order candidate because it provides the best compromise between global accuracy and computational cost, while LE-L9 is retained as a refined CUF reference for local sectional stress recovery.

5. TC2 Benchmark: Composite Box Beam under Coupled Bending and Torsion

TC2 assesses whether the model-order trends observed in TC1 remain valid for a closed composite thin-walled section. The purpose is not to introduce a new hierarchy, but to test the previously defined CUF model families and finite-element reference models on a more representative wing-box-like configuration. The eccentric tip load produces bending-torsion coupling, shear transfer between skins and webs, and local stress redistribution near corner regions. The H8-3 solid model is used as the quantitative three-dimensional reference, while the shell model is retained as an industrial finite-element baseline.

5.1 Global Response of the CUF beam model

Figure 10 reports the relative error in total strain energy as a function of the measured wall-clock solution time for the TC2 benchmark. The graph follows the same plotting logic introduced for the TC1 load cases: each point corresponds to one analysed structural model, the vertical coordinate gives the error with respect to the H8-3 solid reference, and the horizontal coordinate gives the corresponding solution time. The labels associated with the CUF points identify the Taylor-expansion order, so that moving along the TE sequence shows the effect of increasing the global polynomial enrichment of the cross-section. The shell and solid points correspond to the mesh-refinement levels reported in Table 4.

The shell solutions retain a systematic offset from the H8-3 reference even after mesh refinement, whereas the CUF-TE sequence progressively approaches the three-dimensional solution as the Taylor order increases. This behaviour indicates that, for this closed-section benchmark, improving the sectional kinematics through the CUF expansion is more effective for the global bending-torsion response than refining the shell mesh alone. In particular, TE4 already provides a compact approximation with a global error close to 2.6%, while TE8 reduces the strain-energy error below 0.3% at a higher computational cost.

Table 6 – Relative error in total strain energy for TC2 as a function of measured wall-clock solution time. The H8-3 solid model is adopted as the reference solution. Each point corresponds to one analysed model: TE labels indicate Taylor-expansion orders, whereas shell and solid labels indicate the mesh-refinement levels reported in Table 4.

Model	DOFs	Runtime [s]	U	e_U [%]	$u_{z,\max}$	$e_{u_z,\max}$ [%]
H8-1	73440	7.000	12594.484	1.776	25.231	1.613
H8-2	162000	14.000	12697.895	0.969	25.435	0.818
H8-3	544320	54.000	12822.156	-	25.644	-

The solid refinement study keeps the geometry, laminate, boundary conditions, loading, and through-thickness ply discretisation unchanged, while refining only the in-plane mesh of the box walls. From H8-2 to H8-3, the total strain energy changes by approximately 0.98%, while the maximum vertical displacement changes by approximately 0.82%. H8-3 is therefore adopted as the finest available three-dimensional reference for the TC2 global-error assessment. Local peak stresses are considered separately because they remain more sensitive to mesh density and stress recovery in the skin–web corner region.

The quantitative comparison in Table 7 shows that TE8 gives the most accurate global response among the tested Taylor models, with $e_U = 0.249\%$ and $e_{u_z,\max} = 0.119\%$. However, TE4 remains the more attractive model: it uses only 1395 DOFs, runs in about 0.11 s, and keeps both global errors close to 2.6%. By contrast, the refined shell model SH-4 retains errors of about 5.9% and is slower than the H8-3 reference in this benchmark, as indicated by the negative time saving. This result shows that mesh refinement of a shell model does not compensate for missing sectional kinematic richness in the coupled bending-torsion response.

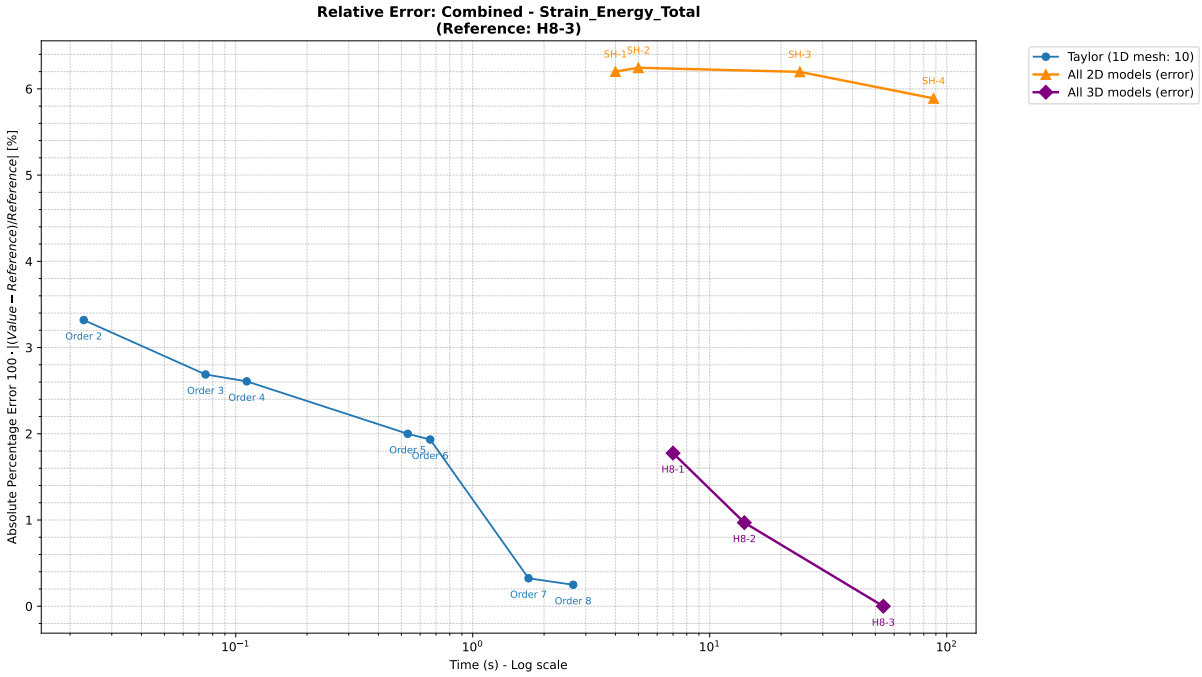


Figure 10 – Total strain-energy error for TC2 as a function of computational time. The H8-3 solid model is adopted as the reference solution.

Table 7 – Summary of TC2 global response, computational cost, and errors with respect to the H8-3 solid reference.

Model	DOFs	Runtime [s]	Time saving [%]	U	e_U [%]	$u_{z,\max}$	$e_{u_z,\max}$ [%]
H8-3 ref.	544320	54.000	-	12822.156	-	25.6444	-
TE1	279	0.023	99.96	11214.447	12.539	22.429	12.539
TE2	558	0.023	99.96	12396.611	3.319	24.793	3.319
TE3	930	0.075	99.86	12477.537	2.688	24.955	2.688
TE4	1395	0.111	99.79	12487.777	2.608	24.976	2.608
TE5	1953	0.532	99.02	12565.760	2.000	25.132	2.000
TE6	2604	0.662	98.77	12574.316	1.933	25.149	1.933
TE7	3348	1.720	96.82	12780.608	0.324	25.581	0.249
TE8	4185	2.653	95.09	12790.265	0.249	25.614	0.119
SH-1	6612	4.000	92.59	12027.146	6.200	24.096	6.037
SH-2	26136	5.000	90.74	12021.450	6.245	24.043	6.245
SH-3	104112	24.000	55.56	12027.574	6.197	24.055	6.197
SH-4	415584	88.000	-62.96	12066.967	5.890	24.134	5.890
H8-1	73440	7.000	87.04	12594.484	1.776	25.231	1.613
H8-2	162000	14.000	74.07	12697.895	0.969	25.435	0.818

The spanwise vertical-displacement error profile in Figure 11 complements the scalar quantities in Table 7. Since no ribs or internal transverse reinforcements are present in TC2, the local peaks are not associated with structural components; they mainly result from the relative-error definition near regions where the reference displacement is small, and from load-introduction and sectional-warping effects near the free end. Over most of the span, TE4 reduces the systematic shell discrepancy, while TE6 and TE8 further improve the agreement with the H8-3 reference.

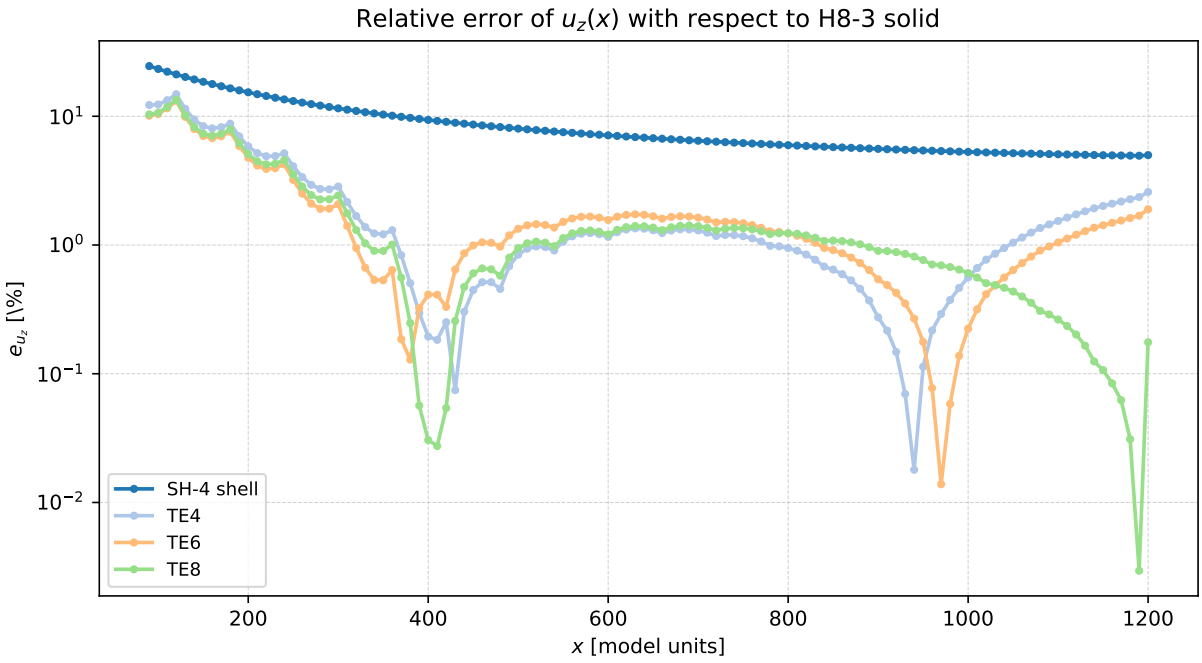


Figure 11 – Relative error of the vertical displacement profile $u_z(x)$ along the beam axis with respect to the H8-3 reference.

5.2 Sectional Response and Model Selection

Global quantities do not fully determine the suitability of a model for composite sizing, since local strain or stress-based constraints may govern the design. For this reason, the TC2 response is also examined at the mid-span section, $x/L = 0.5$, focusing on the lower-left skin–web junction. This region combines bending-induced axial stress, skin–web load transfer, and local stress redistribution near the box corner.

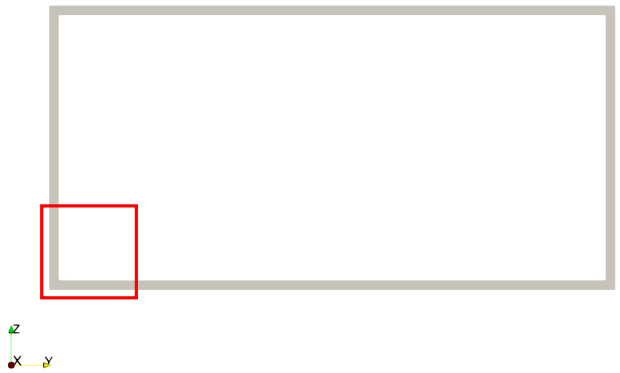


Figure 12 – TC2 mid-span section at $x/L = 0.5$. The highlighted region identifies the lower-left skin–web junction used for the local σ_{xx} comparison.

Figure 13 compares the local non-averaged σ_{xx} field over the highlighted two-dimensional corner region. Stress averaging across ply interfaces or adjacent elements is avoided to preserve laminate-related stress discontinuities. The H8-3 field is used as the three-dimensional reference. TE4 repro-

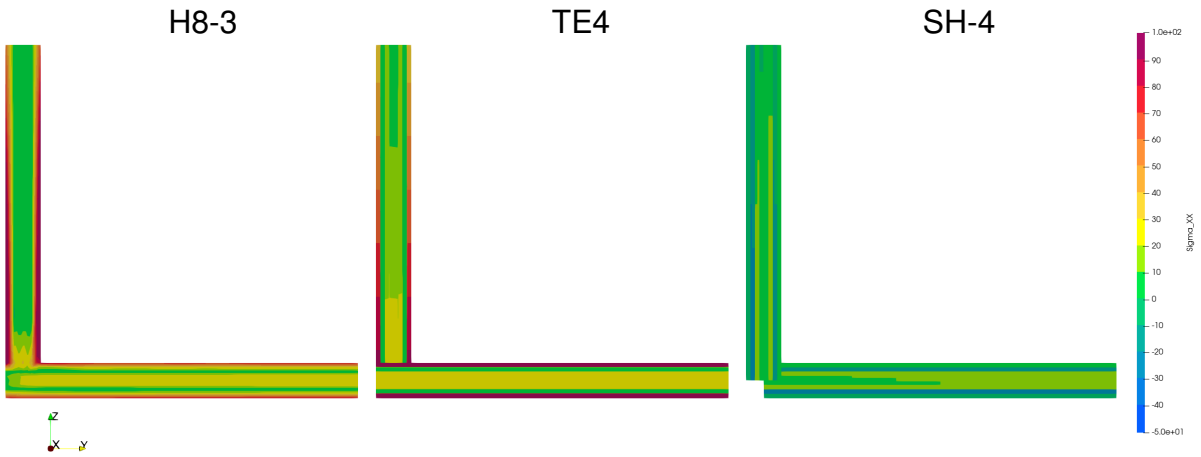


Figure 13 – Local non-averaged σ_{xx} field in the selected lower-left skin–web junction of the TC2 section at $x/L = 0.5$. The same colour scale is used for H8-3, TE4, and SH-4.

duces the main high-stress bands along the skin–web junction, whereas SH-4 provides a smoother field and underestimates the local tensile stress level.

Table 8 – Robust local σ_{xx} comparison in the lower-left skin-web junction of the TC2 section at $x/L = 0.5$. TE4 and SH-4 are interpolated onto the H8-3 sectional points.

Model	Eval. points	$\sigma_{xx,\max}^{\mathrm{raw}}$ [MPa]	σ_{xx}^{P95} [MPa]	$e_{\sigma_{xx}^{P95}}$ [%]
H8-3	402	169.110	86.447	-
TE4	402	99.609	99.572	15.18
SH-4	402	17.320	12.958	85.01

The quantitative comparison in Table 8 is performed over the same two-dimensional corner region. The H8-3 sectional points inside the highlighted region are used as evaluation points, and the TE4 and SH-4 stress fields are interpolated onto those points. Since the raw maximum is affected by isolated local values, the comparison uses the 95th percentile of the local σ_{xx} distribution as a robust indicator of the representative high-stress level:

$$e_{\sigma_{xx}^{P95}} = \frac{\left|\sigma_{xx}^{P95,\mathrm{model}} - \sigma_{xx}^{P95,\mathrm{H8-3}}\right|}{\left|\sigma_{xx}^{P95,\mathrm{H8-3}}\right|} \times 100. \tag{7}$$

Using this metric, TE4 differs from the H8-3 reference by 15.18%, whereas SH-4 underestimates the representative high tensile stress level by 85.01%. TE4 is therefore more suitable than the shell baseline for preliminary local-stress estimation in this corner region. Nevertheless, detailed corner-stress verification should still rely on dedicated three-dimensional models.
Combining the global and local evidence, TE4 is retained as the best trade-off model for the optimisation study. TE8 is more accurate globally, but TE4 provides a better balance between runtime, DOFs, and sufficient sectional fidelity for repeated design evaluations.

5.3 Optimisation Results and Discussion

The preliminary composite sizing problem is formulated as a constrained mass-minimisation problem. The design vector $\mathbf{t} = [t_1, \ldots, t_{N_{\mathrm{DV}}}]^T$ collects $N_{\mathrm{DV}} = 16$ ply-wise thickness variables assigned to the box-beam walls. The stacking sequence and ply orientations are kept fixed, and each design variable is constant along the beam axis, resulting in a spanwise-uniform thickness sizing of the closed section. The variables are treated as continuous equivalent ply-thicknesses, with bounds $t_i^{\min} = 0.125$ mm and $t_i^{\max} = 0.500$ mm, corresponding to one and two nominal ply thicknesses. A final manufacturable design would require subsequent rounding or projection onto discrete ply counts.

The optimisation problem is written as

$$\begin{aligned} \min_{\mathbf{t}} \quad & m(\mathbf{t}) \\ \text{s.t.} \quad & g_k(\mathbf{t}) \leq 0, \qquad k = 1, \ldots, N_g, \\ & t_i^{\min} \leq t_i \leq t_i^{\max}, \qquad i = 1, \ldots, N_{\mathrm{DV}}. \end{aligned} \tag{8}$$

Here, $m(\mathbf{t})$ is the structural mass and $g_k(\mathbf{t})$ denotes the strain-based constraints used in the sizing loop. In the CUF workflow, the TE4 beam model recovers the three-dimensional strain components from the generalized displacement field, and the nodal sectional strain constraints are combined through a differentiable aggregation function compatible with gradient-based SLSQP. The commercial shell-based workflow uses the corresponding shell model with the same number of design variables, but constrains the strain components separately. Since the structural discretisations and constraint treatments are not identical, the two optimisation histories are interpreted as a workflow-level comparison rather than as a strict one-to-one optimiser ranking.
Figure 14 compares the CUF-TE4/SLSQP optimisation with the commercial shell-based SQP workflow. The objective is normalised by the corresponding initial mass because the two analyses use different structural descriptions. The CUF workflow reaches a feasible design in about ten iterations, whereas the shell-based workflow converges more gradually in 25 iterations.

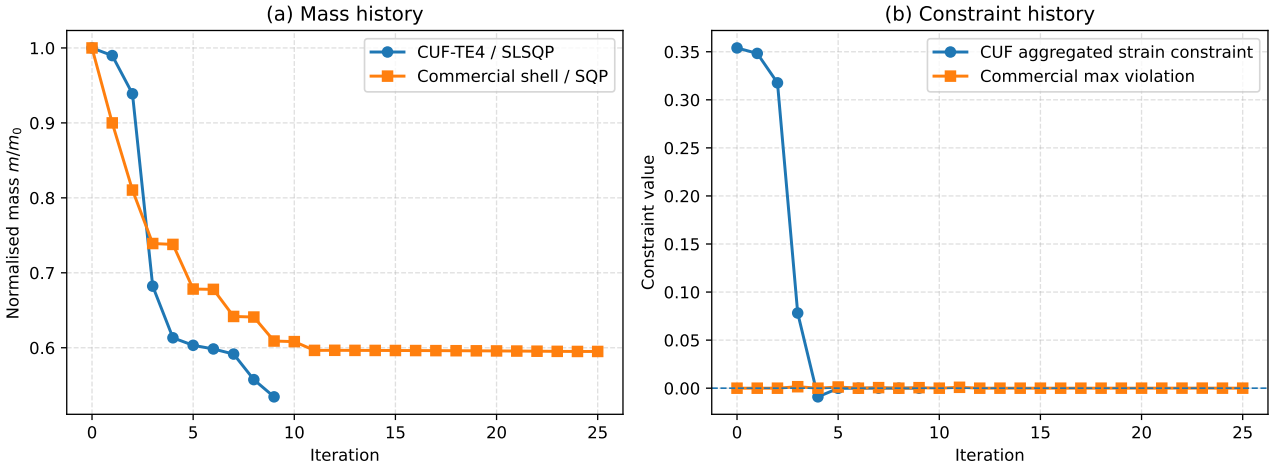


Figure 14 – Optimisation histories for the CUF-TE4/SLSQP and shell-based SQP workflows.

Table 9 – Summary of the CUF-TE4 and commercial shell-based optimisation results. Objective values are reported in the native units of each workflow, while mass reduction is normalised with respect to the corresponding initial value.

Quantity	CUF-TE4/SLSQP	Commercial shell/SQP
Analysis model	TE4 CUF beam	SH-4 shell model
Optimiser	SLSQP	OptiStruct SQP
Sensitivity computation	Adjoint with reverse-mode AD	Internal OptiStruct sensitivities
Number of design variables	16	16
Initial objective [kg]	1.351680	1.38240
Final objective [kg]	0.722435	0.822591
Normalised mass reduction [%]	46.55	40.50
Number of iterations	10	25
Final constraint status	Feasible, active constraint	Feasible, soft convergence

The CUF-TE4 workflow reaches a feasible design in 10 iterations and achieves a 46.55% normalised mass reduction, whereas the shell-based workflow converges more gradually and reaches a 40.50% reduction. These values should not be interpreted as a direct ranking of the optimisers, because the constraint aggregation and structural discretisations differ. Instead, the result demonstrates that the selected TE4 model can be effectively embedded in a preliminary sizing loop while retaining a very low analysis cost.

Figure 15 shows the final thickness distributions obtained from both sizing workflows. The resulting layouts are mechanically consistent with the coupled bending–torsion response. In the CUF solution, most variables approach the lower bound, while material is retained mainly in selected 0° plies of skins and spars. The shell-based optimisation favours more material in the front spar and selected skin plies. Overall, the optimisation confirms that TE4 is suitable for preliminary sizing loops, provided that the final design is subsequently verified with higher-fidelity models and consistent constraint definitions.

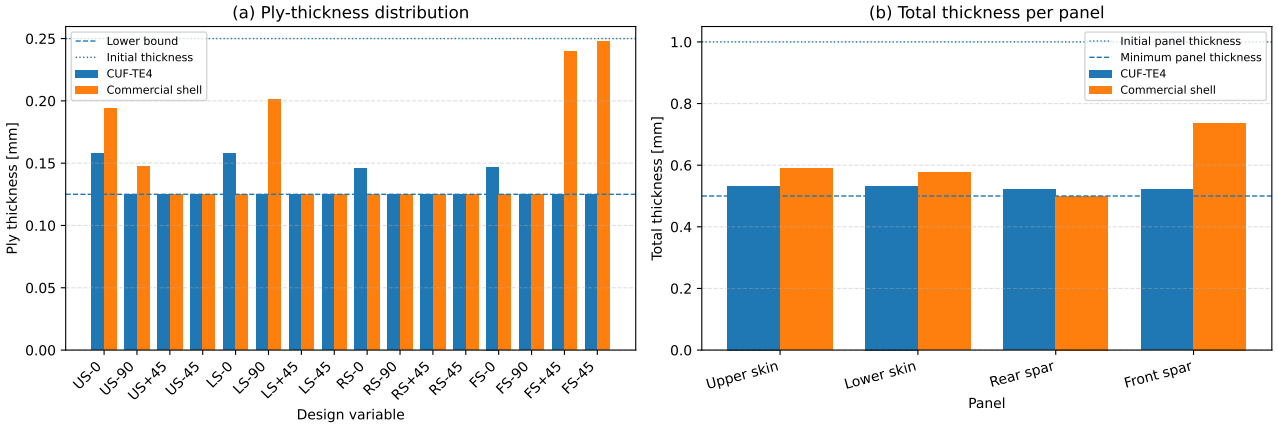


Figure 15 – Final ply-wise and panel-wise thickness distributions.

The present optimisation should therefore be interpreted as a preliminary sizing exercise rather than a certification-level design procedure. Buckling, detailed composite failure criteria, manufacturing constraints, stiffeners, ribs, cut-outs, and load-introduction effects are not included at this stage. These aspects would need to be considered in a subsequent high-fidelity verification and detailed sizing phase.

6. Conclusions

This work assessed hierarchical CUF beam models for optimisation-oriented analysis of laminated composite wing-like structures. Taylor- and Lagrange-expansion CUF models were compared with shell and solid finite-element references using two benchmarks: a laminated flat panel under elementary axial, bending, and torsional loads, and a closed composite box beam under coupled bending and torsion.

For the TC1 flat-panel benchmark, the results showed that the required sectional enrichment is mechanism-dependent. Axial loading is accurately captured by low-order models, bending benefits from moderate enrichment, and torsion is the most discriminating case. Across the three load cases, TE4 kept the strain-energy error below 0.5% while reducing the measured solution time by more than 99.8% with respect to the solid reference.

For the TC2 box-beam benchmark, TE8 provided the most accurate global response among the tested Taylor models, with errors below 0.3%. However, TE4 offered the best accuracy–cost compromise, with global errors close to 2.6%, 1395 degrees of freedom, and a runtime of approximately 0.11 s. The local σ_{xx} comparison in the skin–web junction also showed that TE4 gives a more representative preliminary estimate of the high-stress region than the refined shell baseline, although detailed corner-stress verification still requires dedicated three-dimensional models.

The selected TE4 model was then embedded in a preliminary thickness-sizing loop, achieving a 46.55% normalised mass reduction in 10 iterations. The optimisation results should be interpreted as a workflow-level demonstration rather than a strict one-to-one comparison with the shell-based procedure, because the structural discretisations and constraint treatments differ. Overall, the results indicate that moderate-order CUF beam models can provide an effective modelling route for early-stage composite wing-box sizing, with higher-order CUF and solid models retained for verification and final local-stress assessment.

Acknowledgements

This work was conducted within the project TIFON (Tecnologías Inteligentes para la Fabricación, el Diseño y las Operaciones en entornos Industriales), coordinated by the Universidad Politécnica de Madrid (UPM).

7. Contact Author Email Address

dcardone@ign.uc3m.es

8. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] J. N. Reddy. *Mechanics of Laminated Composite Plates and Shells: Theory and Analysis*. CRC Press, 2 edition, 2004.

[2] László P. Kollár and George S. Springer. *Mechanics of Composite Structures*. Cambridge University Press, 2003.

[3] Liviu Librescu and Ohseop Song. *Thin-Walled Composite Beams: Theory and Application*. Springer, 2006.

[4] Thuc Vo and Jaehong Lee. Flexural-torsional behavior of thin-walled closed-section composite box beams. *Engineering Structures*, 29(8):1774–1782, 2007.

[5] Thuc Phuong Vo and Jaehong Lee. Flexural-torsional behavior of thin-walled composite box beams using shear-deformable beam theory. *Engineering Structures*, 30(7):1958–1968, 2008.

[6] Chuandi Zhang and Shijun Wang. Structure mechanical modeling of thin-walled closed-section composite beams, part 1: Single-cell cross section. *Composite Structures*, 113:12–22, 2014.

[7] Shijun Wang and Chuandi Zhang. Structure mechanical modeling of thin-walled closed-section composite beams, part 2: Multi-cell cross section. *Composite Structures*, 113:56–62, 2014.

[8] François-Xavier Irisarri, Frédéric Laurin, François-Henri Leroy, and Jean-François Maire. Computational strategy for multiobjective optimization of composite stiffened panels. *Composite Structures*, 93(3):1158–1167, 2011.

[9] Sridhar Chintapalli, Mostafa S. A. Elsayed, Ramin Sedaghati, and Mohammed Abdo. The development of a preliminary structural design optimization method of an aircraft wing-box skin-stringer panels. *Aerospace Science and Technology*, 14(3):188–198, 2010.

[10] Zexi Wang, Zhiqiang Wan, Rainer M. J. Groh, and Xiaozhe Wang. Aeroelastic and local buckling optimisation of a variable-angle-tow composite wing-box structure. *Composite Structures*, 258:113201, 2021.

[11] Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Cambridge University Press, 2021.

[12] Erasmo Carrera and Gaetano Giunta. Refined beam theories based on a unified formulation. *International Journal of Applied Mechanics*, 2(1):117–143, 2010.

[13] Erasmo Carrera and Marco Petrolo. Refined one-dimensional formulations for laminated structure analysis. *AIAA Journal*, 50(1):176–189, 2012.

[14] Erasmo Carrera and Alfonso Pagani. Carrera unified formulation for free-vibration analysis of aircraft structures. *AIAA Journal*, 54(1):280–292, 2016.

[15] Marco Petrolo and Erasmo Carrera. Best theory diagrams for multilayered structures via shell finite elements. *Advanced Modeling and Simulation in Engineering Sciences*, 6:4, 2019.

[16] Erasmo Carrera and Alfonso Pagani. Evaluation of the accuracy of classical beam fe models via locking-free hierarchically refined elements. *International Journal of Mechanical Sciences*, 100:169–179, 2015.

[17] Erasmo Carrera, Alfonso Pagani, Marco Petrolo, and Enrico Zappino. Recent developments on refined theories for beams with applications. *Mechanical Engineering Reviews*, 2(2):14–00298, 2015.

[18] Donato Cardone, Rauno Cavallaro, Andrea Cini, Marco Petrolo, and Enrico Zappino. Application of CUF for the structural optimization of wings. In *AIAA SCITECH 2025 Forum*, January 2025.

[19] Donato Cardone, Rauno Cavallaro, and Andrea Cini. Gradient-based optimisation of composite aircraft structures using high-order beam models. *Engineering Proceedings*, 133(1):98, 2026.