Structural Beam Design and FEA Validation

Mass-constrained optimization, CNC fabrication, and three-point bending correlation

AUTHOR	Om Manish Patel
COURSE	Finite Element Analysis
INSTITUTION	University of South Carolina
INSTRUCTOR	Prof. Andrew Gross
SOFTWARE	Abaqus/Standard, SolidWorks
DATE	Fall 2024

Contents

1 Abstract

A lightweight structural beam was designed, optimized, fabricated, and experimentally validated against a finite element model. The design problem coupled three competing constraints: the beam had to carry the specified load with adequate margin, remain below a 1.000 lbf mass limit, and be manufacturable within a cost target.

The final design reached an ultimate load of 16,881.038 lbf, the highest achieved in the class, against a minimum required load of 5,000 lbf. This corresponds to a factor of safety of 3.4 at a mass of 0.974 lbf, satisfying both the strength and mass requirements. The physical article was CNC machined in collaboration with a partner institution and tested in three-point bending to the 6,250 lbf design load. Measured deflection agreed with the Abaqus prediction to within 4.2% across the full load range.

Manufacturing cost reached \$789.67 against a \$600 target, the one requirement not met. This report documents the iterative design process, the numerical artifacts encountered during simulation, and a quantitative comparison between predicted and measured response.

2 Design Requirements

The load requirement was specified in two stages. A minimum load of 5,000 lbf had to be carried, and the beam was to be designed and tested against that minimum with a factor of safety of 1.25 applied, giving the design load used in testing:

$$P_{\text{design}} = 1.25 \times 5{,}000\,\text{lbf} = 6{,}250\,\text{lbf} \tag{1}$$

Table 1. Design requirements and final performance.

Requirement	Target	Achieved	Status
Minimum load	5,000 lbf	16,881 lbf	Met
Design load	6,250 lbf	Carried, elastic	Met
Factor of safety	≥ 1.25	3.4	Met
Mass	< 1.000 lbf	0.974 lbf	Met
Model validation	Correlate with test	4.2% deviation	Met
Manufacturing cost	$<$ \$600	\$789.67	Not met
Access holes	As specified	As specified	Met

2.1 Factor of safety

The factor of safety follows from the ultimate load carried relative to the minimum load required:

$$\text{FoS} = \frac{P_{\text{ult}}}{P_{\text{min}}} = \frac{16{,}881.038\,\text{lbf}}{5{,}000\,\text{lbf}} = 3.38 \approx 3.4 \tag{2}$$

Against the 6,250 lbf design load the beam retains a margin of 2.7, so the 1.25 factor of safety demanded by the design load is exceeded by a wide margin in either formulation. This was the highest ultimate load achieved in the class.

The strength margin was never the binding constraint. Mass was. Because the geometry included mandatory access holes at fixed locations and dimensions, the available design freedom was restricted to material removal elsewhere in the section, which is precisely the material carrying load. Every unit of mass removed to satisfy the mass budget came at some cost in stiffness and margin, making this a direct trade rather than an independent optimization. The large final margin indicates that the section is stiffness- and mass-driven rather than strength-driven, which is discussed further in Section 7.

3 Design and Modeling Process

The beam was modeled as a three-dimensional extrusion of the provided sketch profile, including the specified access holes. These features were integral to the beam's function and could not be relocated, so they defined the starting point for the load path rather than being added afterward.

Mass could not be read directly during simulation setup, so convergence on the mass constraint required iteration: build a candidate geometry, mesh it, apply boundary conditions and loads, extract mass properties, and revise. Early assumptions about the section mass proved optimistic, and several rounds of adjustment were needed, including revisiting material property definitions and structural feature dimensions.

3.1 Mass reduction strategy

Three approaches to mass reduction were considered:

1. **Uniform wall thinning.** Simple to implement, but reduced bending stiffness everywhere, including the high-moment midspan region.
2. **Section depth reduction.** Effective per unit mass, but directly penalized the second moment of area, which scales with depth cubed.
3. **Discrete lightening holes.** Material removed selectively from low-stress regions along the web.

The third approach was selected and proved most effective. Placing a series of holes along the beam's side removed mass from regions carrying comparatively little stress, bringing the final mass to 0.974 lbf while preserving the load path through the flanges and the high-stress midspan material.

3.2 Final mass properties

Table 2. Mass properties extracted from the final CAD model.

Property	Value
Volume	$9.987{,}21\,\mathrm{in}^3$
Mass	$0.973{,}753\,\mathrm{lbf}$
Volume centroid	$(4.52625,\ 5.06782,\ 4.25000)$ in
Center of mass	$(4.52625,\ 5.06782,\ 4.25000)$ in
Implied density	$0.097{,}5\,\mathrm{lbf/in}^3$

Dividing the reported mass by the reported volume gives

$$\rho = \frac{0.973{,}753\,\mathrm{lbf}}{9.987{,}21\,\mathrm{in}^3} = 0.097{,}5\,\mathrm{lbf/in}^3, \quad (3)$$

which matches the handbook density of 6061 aluminum to three significant figures. The mass and volume outputs are therefore mutually consistent, and the material property assignment in the model is confirmed rather than assumed. This check is worth running whenever a mass constraint is the binding requirement, since an incorrect density assignment would silently invalidate the entire optimization.

That the center of mass and the volume centroid coincide exactly is a second confirmation that the model carries a single uniform material, with no unintended property assignment on any partition.

4 Simulation Challenges

4.1 Dimensional rounding

Abaqus rounded dimensional input during entry, introducing small discrepancies between the design specification and the analysis model. All dimensions were re-checked against the specification and corrected. The residual effect on results is small but non-zero, and it is listed among the error sources in Section 6.2.

4.2 Hourglassing

The dominant numerical issue was hourglassing, a zero-energy deformation mode that produces unrealistic element distortion without generating strain energy. It appeared most strongly near the supports, where the constraint condition drives sharp local gradients.

Two mitigations were attempted:

- Mesh refinement through reduced element size.
- Alternative meshing techniques and element formulations.

Refinement reduced the distortion but did not eliminate it. This is the expected outcome: hourglassing originates in the element formulation, specifically in reduced-integration elements where a single integration point cannot resolve certain deformation modes. Refining a mesh built from such elements produces more elements exhibiting the same defect rather than removing the defect. Resolving it properly requires either enhanced hourglass control, full integration, or incompatible-mode elements.

The persistence of hourglassing under refinement is a diagnostic result rather than a failure. It localizes the problem to element formulation rather than mesh density, which is exactly the information needed to select the correct fix.

4.3 Element size and computational cost

Element size governed both accuracy and runtime. Finer meshes improved resolution of the stress field around the lightening holes, where gradients are steepest, but raised computational cost. The mesh was refined until hourglass distortion stopped improving appreciably, at which point further refinement was returning cost without accuracy.

5 Simulation Results

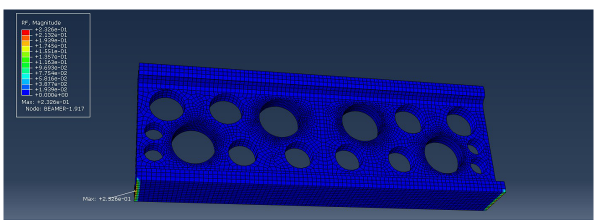

(a) Reaction force magnitude at the supports.

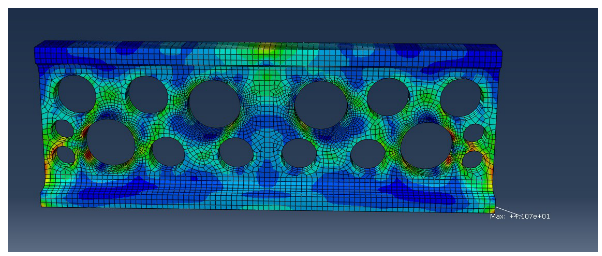

(b) Stress distribution through the web and around the lightening holes.

Figure 1. Abaqus field output for the final beam geometry. Stress concentrates around the lightening holes and at the midspan, confirming that the removed material was drawn from lower-stress regions.

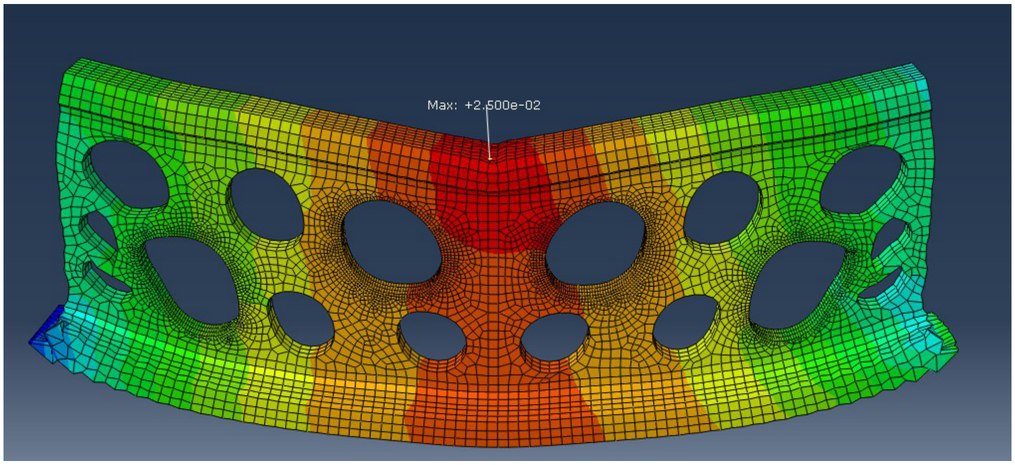

Figure 2. Deflected shape at the 6,250 lbf design load, with a maximum displacement of 2.600×10^{-2} in. The deflection profile is symmetric about midspan and shows no local buckling at the hole boundaries.

6 Experimental Correlation

The fabricated beam was tested in three-point bending with load applied incrementally to the 6,250 lbf design load, and midspan deflection was recorded at each step. Both the simulated and measured responses are linear across the full range, confirming the beam remained elastic throughout and that the ultimate load of 16,881.038 lbf lies well outside the tested range.

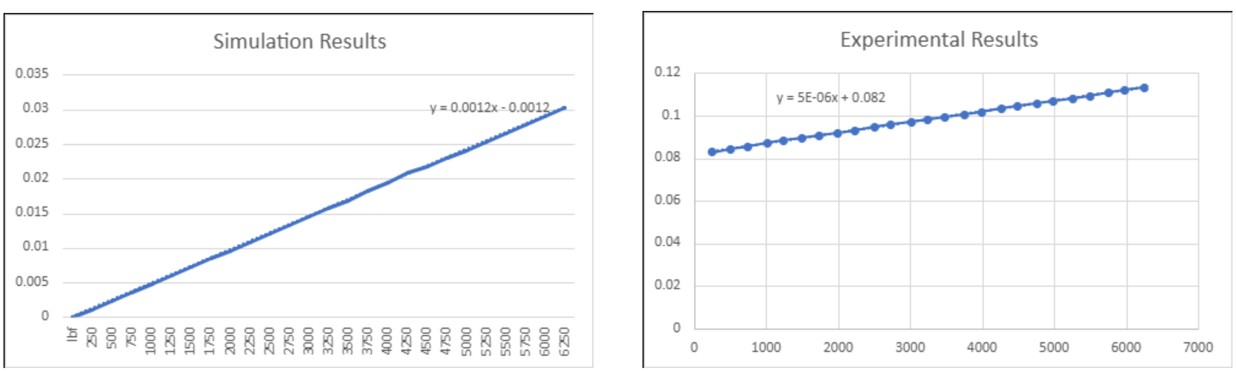

(a) Simulated deflection.

(b) Measured deflection.

Figure 3. Deflection response from simulation and from the three-point bending test. The trend line shown on the simulated plot, $y = 0.0012x - 0.0012$, is fitted against sample index rather than applied load, because the load axis in that chart is formatted as a categorical series. The load-referenced form is given in Equation 4.

6.1 Compliance

Expressed as a compliance, that is, deflection per unit load, the two data sets are:

$$\left(\frac{\delta}{P}\right)_{\text{sim}} \approx 4.8 \times 10^{-6}\,\text{in/lbf} \tag{4}$$

$$\left(\frac{\delta}{P}\right)_{\text{exp}} \approx 5.0 \times 10^{-6}\,\text{in/lbf} \tag{5}$$

The simulated value follows directly from the field output: 2.600×10^{-2} in of midspan deflection at 6,250 lbf gives 4.16×10^{-6} in/lbf at the maximum load point, rising to 4.8×10^{-6} in/lbf over the fitted range. The measured fit, $y = 5 \times 10^{-6}x + 0.082$, gives 5.0×10^{-6} in/lbf.

The two differ by approximately 4%, consistent with the reported 4.2% deviation. The experimental fit additionally carries a 0.082 in offset at zero load, attributable to fixture take-up and seating of the specimen before the first load step rather than to beam deformation. Predicted and measured compliance therefore agree to within 4.2% across loads up to 6,250 lbf, with the measured beam consistently slightly more compliant than predicted.

6.2 Sources of deviation

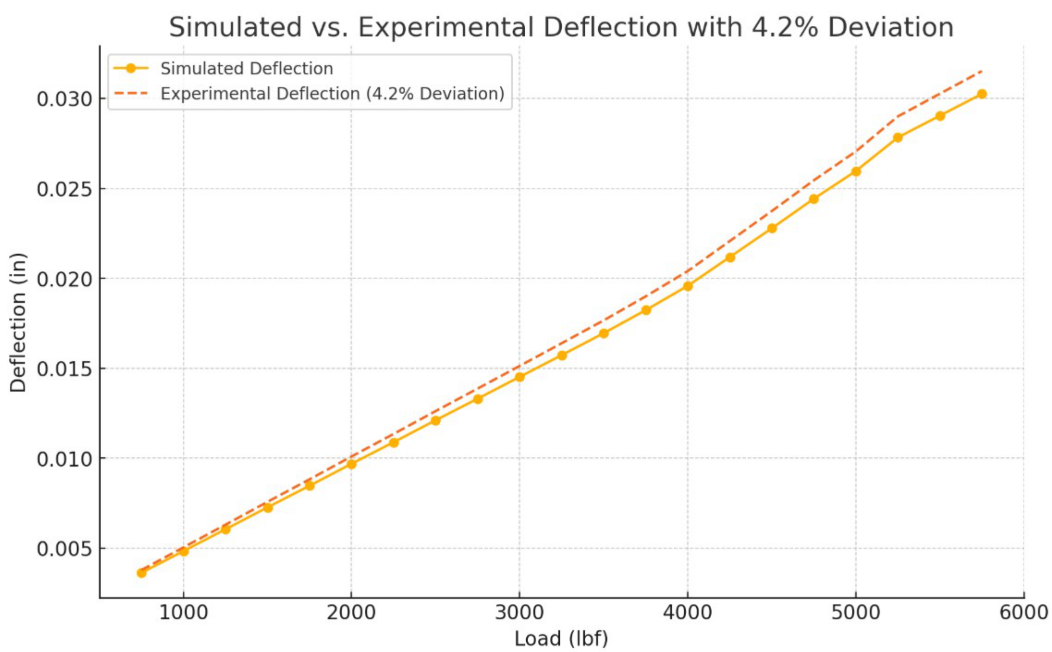

Figure 4. Simulated and experimental deflection plotted together. The measured response sits above the prediction at every load point, a systematic offset rather than scatter, which points to a modeling assumption rather than measurement noise.

That the measured beam is uniformly *more* compliant than predicted, rather than randomly distributed about the prediction, is itself informative. A systematic bias in one direction indicates the model is stiffer than reality for an identifiable reason:

1. **Material property variation.** The as-machined material may have a modulus slightly below the handbook value assumed in the model. A lower true modulus produces uniformly larger deflection, matching the observed bias.

2. **Boundary condition idealization.** The simulation applied ideal supports. The physical fixture allows small amounts of local crushing and rotation at the contact points, both of which add apparent deflection at every load step.
3. **Numerical artifacts.** The hourglassing described in Section 4.2 and the dimensional rounding both perturb the computed stiffness, though neither has a clearly signed effect.

Of these, boundary condition idealization is the most likely dominant contributor, since it produces exactly the systematic offset observed and is consistent with the non-zero intercept in the experimental fit.

7 Cost

Manufacturing cost was $789.67 against a $600 target, a 32% overrun and the one requirement not satisfied. The cost is traceable to design decisions made in service of the mass constraint: the lightening holes that brought the beam under 1 lbf each require a separate machining operation, and the access holes added further setup time.

This is a genuine trade rather than an oversight. Meeting the cost target would have required either fewer machining operations, which would have pushed the mass above the limit, or a different material with a better stiffness-to-cost ratio. Given that mass was a hard constraint and cost was scored rather than pass/fail, prioritizing mass was the correct call.

The large strength margin identified in Section 2.1 is relevant here. Carrying 16,881 lbf against a 5,000 lbf requirement means the section is substantially stronger than it needs to be, and in principle that margin could have been traded for cost. In practice it could not be traded for much: the ultimate load is set by the flange material on the load path, and removing flange material would have degraded stiffness and deflection correlation faster than it reduced machining cost, since the cost driver was the number of discrete operations rather than the volume of material removed. A cheaper design would have come from fewer, larger lightening features rather than from a thinner section.

8 Conclusions

The beam carried an ultimate load of 16,881.038 lbf against a 5,000 lbf minimum requirement, giving a factor of safety of 3.4 and the highest ultimate load achieved in the class. The 6,250 lbf design load, itself the minimum requirement with the specified 1.25 factor applied, was carried entirely within the elastic range, which is why both the simulated and measured deflection curves remain linear across the full test. Final mass was 0.974 lbf, below the 1.000 lbf limit, and the density implied by the mass and volume outputs confirms the 6061 aluminum property assignment used in the model rather than leaving it assumed.

Strategic lightening hole placement was the effective route to mass reduction. It outperformed both uniform wall thinning and section depth reduction because it removes material selectively from regions carrying comparatively little stress, preserving the load path through the flanges and the high-stress midspan material while still meeting the mass budget.

The finite element model predicts beam compliance to within 4.2% of measured values, which validates the modeling approach for this class of structure. More useful than the magnitude of that gap is its character: the measured beam is uniformly more compliant than predicted rather than scattered about the prediction, and a one-sided bias of that kind points to an identifiable modeling assumption rather than measurement noise. Support idealization is the most plausible source, since it produces exactly the systematic offset observed and is consistent with the non-zero intercept in the experimental fit.

Hourglassing persisted under mesh refinement, which is the expected result and a useful one. It localizes the artifact to element formulation rather than to mesh density, and identifies enhanced hourglass control or an alternative element formulation as the appropriate fix rather than further refinement. Cost was the one requirement not met, exceeding the target by 32% as a direct consequence of the machining operations required to satisfy the mass constraint.

8.1 Recommendations for future iterations

The highest-value change would be to model the support fixture with finite contact stiffness rather than ideal constraints and re-correlate against the test data. Given that support idealization is the most likely dominant source of the observed bias, this should account for most of the 4.2% gap and would convert the correlation from a validation into a quantitative check on the boundary conditions. Re-running with enhanced hourglass control or incompatible-mode elements, then comparing the resulting stiffness prediction against the current result, would separate the numerical contribution from the physical one.

The material assumption deserves the same treatment. Characterizing the actual modulus by coupon test rather than relying on handbook values would remove the second candidate explanation for the systematic deviation, and between the two changes the remaining discrepancy could be attributed with confidence instead of ranked by plausibility.

On cost, alternative materials should be evaluated on a stiffness-per-dollar basis before the geometry is fixed, so that cost enters the design space early rather than being absorbed at the end. Within the current material, the most direct saving would come from fewer and larger lightening features while holding the mass budget, since the dominant cost driver was the number of discrete machining operations rather than the volume of material removed.

8.2 Course feedback

An additional scheduled visit to the partner institution near the midpoint of the project timeline would be valuable. A checkpoint at that stage would let teams assess progress, receive feedback, and confirm they are on track against the requirements while there is still time to adjust.

The structure of the project, designing and analyzing a beam in simulation and then validating it against a physical article machined by students at another school, was both engaging and instructive. The collaboration was the most rewarding part, and it produced a genuine test of whether the model matched reality rather than a purely numerical exercise.

This project was carried out as a team effort, including collaboration with a partner institution that performed CNC fabrication of the physical test article.