iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: VII Month of publication: July 2026

DOI: https://doi.org/10.22214/ijraset.2026.84267

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

AI-Assisted Extension of a Nonlinear FEA Program to Material Plasticity: Benchmark-Driven Validation against Commercial FEA

Sreekanth Manakkattil Sivaraman
Indian Institute of Technology Bombay, Subsea Structural Engineer, Malaysia

Abstract: This paper extends a previously published AI-assisted nonlinear finite element analysis library for large-displacement beam structures to material plasticity, under the same human-AI collaborative workflow used for the geometrically-nonlinear original. Two constitutive paths are implemented: a Ramberg-Osgood deformation model and a J2 incremental plasticity model with isotropic hardening, fibre section integration, and a Newton radial-return algorithm required for tabular hardening curves. The implementation is validated against Abaqus B31 beam elements and, decisively, against an independent S4R shell reference model that carries no beam-section idealisation. The converged fibre-integration section tracks the shell reference to within 4-8% across the full load range, a level consistent with ordinary beam-versus-shell theory. A circumferential point-count study of the Abaqus PIPE section (8, 16, 64 points) shows monotonic convergence toward the same fibre answer, while the default PIPE section diverges increasingly from the shell reference as plasticity deepens, reaching about 65% in tip rotation near the fully-plastic moment. Across a multi-support displacement-controlled bend matrix and a moving-support passage test at moderate plasticity, the library agrees with Abaqus to single-digit percentages, and the incremental model reproduces the freezing of plastic strain under elastic unloading confirmed independently in the commercial code. The paper contributes a verified open plasticity extension, a shell-validated characterisation of the commercial PIPE section's plastic-range behaviour, practical guidance on material-definition and validation-metric pitfalls, and an extended taxonomy of errors specific to AI-assisted development of numerical software.
Keywords: Nonlinear finite element analysis, J2 plasticity, incremental plasticity, Ramberg-Osgood, beam elements, section integration, shell model, benchmark verification, large language model, AI-assisted software development, offshore pipeline

I. INTRODUCTION

The companion paper [1] documented a structured human-AI collaborative workflow in which the large language model Claude (Anthropic) was guided by the author to reimplement, verify, and optimise a nonlinear finite element analysis (FEA) program for large-displacement 2D beam structures, originally published in 2015 [2]. That work covered geometric nonlinearity only, verified against NAFEMS benchmarks with linear elastic material. Section VIII of [1] identified material nonlinearity via the Ramberg-Osgood model as ongoing work. The present paper delivers that extension and goes substantially beyond it: alongside the Ramberg-Osgood deformation path, a full J2 incremental plasticity model with isotropic hardening is implemented, enabling the load-unload-reload analysis that deformation plasticity cannot represent.

The motivation is direct. During S-Lay pipeline installation every material point of the pipeline experiences a complete loading cycle as it travels over the stinger rollers: bending builds as the point approaches a roller, peaks at the roller, and unloads as the point moves on, leaving residual plastic strain that a path-independent model discards. Accurate accumulated-strain prediction therefore requires incremental plasticity.

The central technical contribution emerged from the verification campaign. Benchmarking the new implementation against Abaqus B31 beam elements initially showed the converged fibre-integration model predicting less tip rotation than the default Abaqus PIPE section, with the gap growing with load level. Rather than tuning the discrepancy away, a structured investigation was conducted, and, importantly, an independent reference free of any beam-section idealisation, an S4R shell model, was introduced as the arbiter. The result reverses the naive reading: it is the converged fibre model, not the default beam section, that tracks the shell reference; the default Abaqus PIPE section systematically under-integrates the plastic bending response, diverging from physical reality as plasticity deepens. Section VI documents the investigation, its several false turns, and its resolution.

Four contributions are made. First, a verified, open J2 incremental plasticity extension of the library published with [1], with fibre section integration and a Newton radial-return algorithm for tabular hardening curves. Second, a shell-validated characterisation of the plastic-range behaviour of the Abaqus PIPE beam section, including an interchangeable section model that reproduces its coarse-quadrature behaviour for compatibility. Third, practical guidance on two pitfalls that dominated the investigation: silent material-definition mismatches, and the amplification of small section-level differences when comparison is made at matched load rather than matched curvature. Fourth, an extension of the error taxonomy of [1] with error classes specific to AI-assisted development of plasticity algorithms.

II. BACKGROUND

A. Deformation versus incremental plasticity for pipe bending

Deformation (total-strain) plasticity treats stress as a single-valued function of current total strain. The Ramberg-Osgood form [3] is standard in pipeline practice and underlies strain-based design in DNV-ST-F101 [4]. It is adequate for monotonic loading but has no memory: on unloading it retraces the loading curve, predicting zero residual strain. Incremental (flow) plasticity [5] decomposes strain into elastic and plastic parts and evolves the plastic part through a yield condition and flow rule, capturing elastic unloading, residual strain, and reloading. For a material point passing over stinger rollers, the two theories diverge exactly where it matters.

B. Section integration in beam plasticity

Beam elements treat material nonlinearity by numerically integrating stress over the cross-section. For circular hollow sections two idealisations are common: through-thickness fibre integration, discretising the true wall from inner to outer radius, and the thin-wall approximation, lumping all wall material on the mid-wall ring and integrating only around the circumference. The thin-wall elastic error is small for typical pipeline D/t (0.30% on the second moment of area for the 16-inch by 21 mm section used here). Its plastic-range accuracy depends on the number and placement of integration points, which is the crux of Section VI.

C. Reference models

Two independent references are used. Abaqus B31 is a 2-node, shear-flexible 3D beam element with one length-wise integration point; with a PIPE profile the section uses the thin-wall idealisation. An S4R shell model of the same pipe carries no beam-section idealisation at all: it resolves the wall stress distribution directly and therefore serves as a physics reference against which every beam idealisation, the present library's and Abaqus's alike, can be judged. All Abaqus runs used Abaqus/Standard 2024 with NLGEOM active and an identical tabular plasticity definition, so material data cannot be a source of code-to-code difference.

III.FORMULATION

A. Geometric framework (unchanged from [1])

The Combined Corotational-Total Lagrangian formulation of [1], [2] is carried over unchanged: a corotational frame separating rigid-body rotation from deformation, De Souza continuous chord-angle tracking, geometric stiffness including axial and moment terms, and load-controlled Newton-Raphson with adaptive increments. It passed NAFEMS NLGB2/NLGB4 in [1]; a dedicated elastic-only run additionally confirms exact agreement with elastica theory to 0.5 rad tip rotation.

B. Ramberg-Osgood deformation path

The R-O path evaluates the section by fibre integration of the total-strain stress function, inverted numerically at each fibre. A conversion routine builds an equivalent incremental hardening table from any R-O curve for consistency on monotonic loading.

C. J2 incremental plasticity with Newton radial return

Each fibre carries committed plastic strain and accumulated equivalent plastic strain, stored per element and Gauss point and committed only on increment convergence so that iterations within an increment cannot corrupt history. For uniaxial fibre stress the radial return reduces to a scalar problem: find the plastic multiplier such that the returned stress lies on the current yield surface. The common one-step solution is exact only for constant hardening modulus. For the tabular material used here the hardening modulus falls from about 65 GPa near first yield to about 7 GPa at large strain; the one-step formula then under-corrects, leaving stress above the yield surface. The implementation therefore iterates the return with a Newton loop (30 iterations maximum, 1 Pa tolerance on the consistency condition), converging in one step for linear hardening and two to three for the tabular curve.

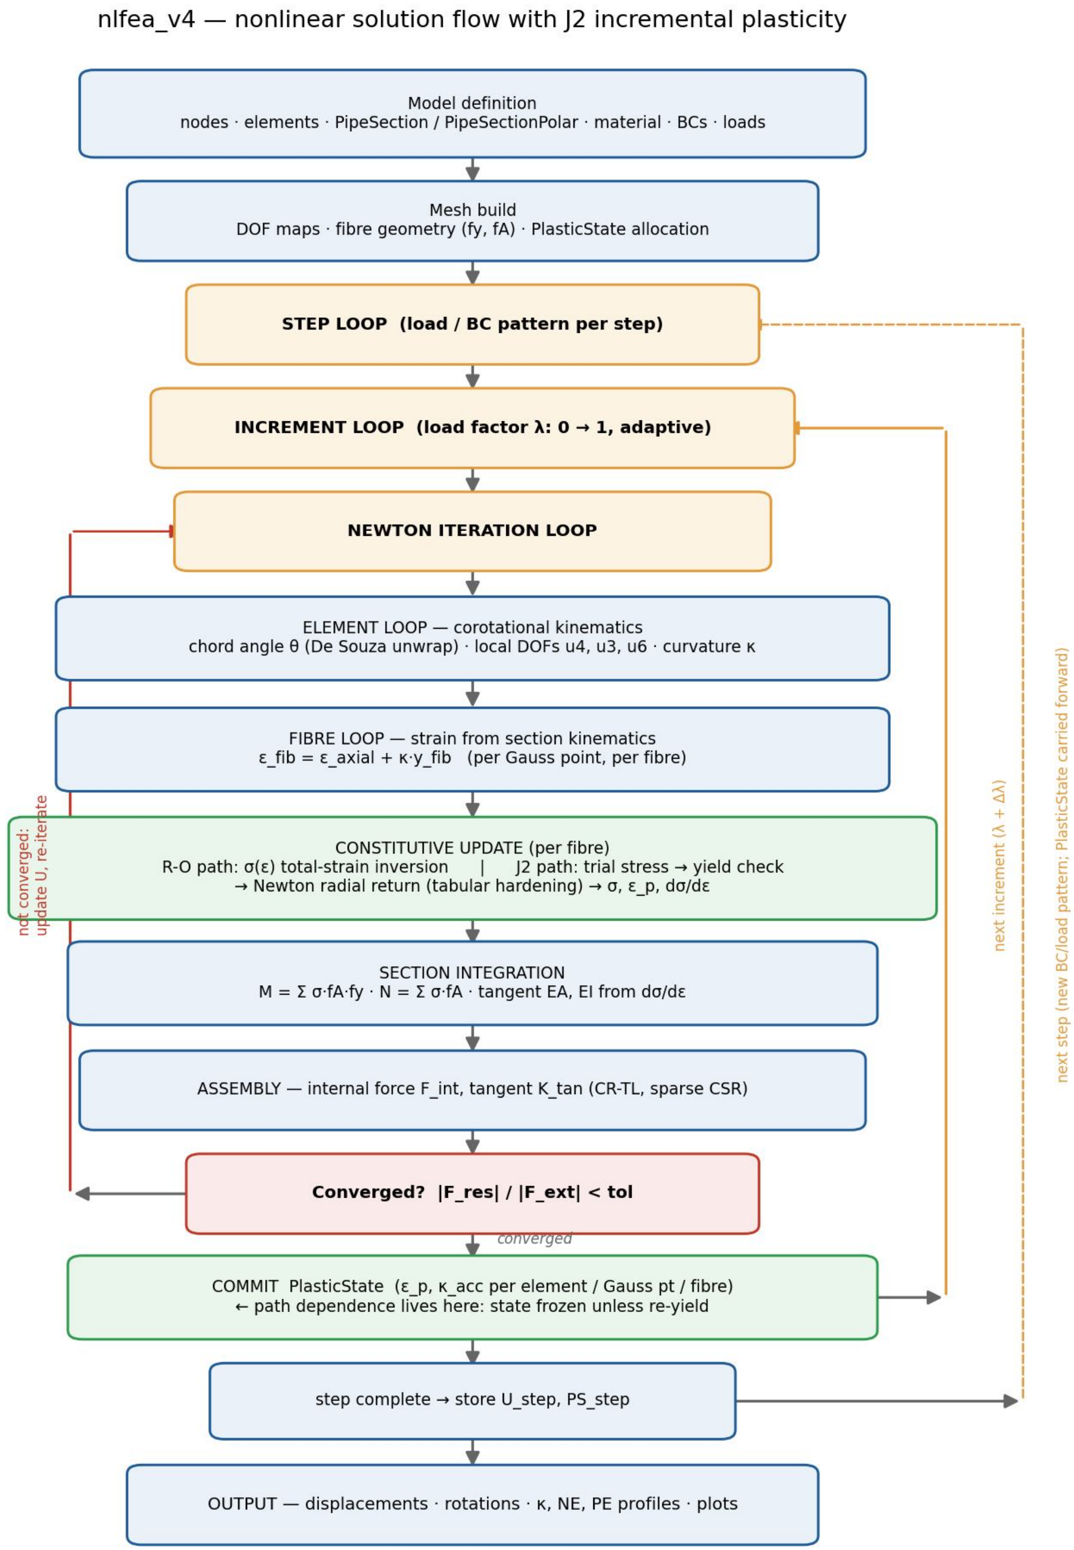

Fig. 1. Solution flow of the library: nested step / increment / Newton loops; the constitutive update and the committed PlasticState (highlighted green) carry the path dependence.

D. Section models

Two interchangeable circular-hollow-section models are provided. The fibre model (denoted exact) discretises the true wall with N horizontal strips across the diameter; a convergence study (Section VI.B) shows N = 200 converged to better than 0.01%. The polar model replicates the Abaqus thin-wall PIPE idealisation, placing all wall area on the mid-wall ring at N equal-area angular points. The default N = 8 integrates the elastic stiffness exactly but under-integrates the fully plastic modulus, a property examined against data in Section VI.

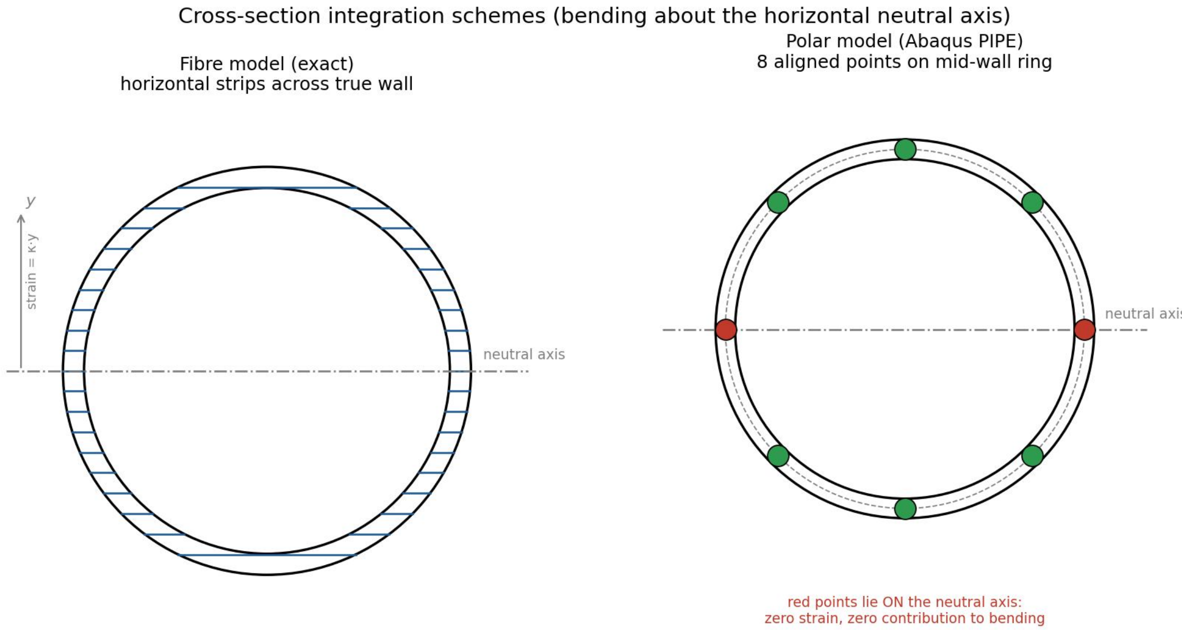

Fig. 2. Cross-section integration schemes, bending about the horizontal neutral axis: fibre strips across the true wall (left) versus 8 aligned points on the mid-wall ring (right). The two points lying on the neutral axis carry zero strain and contribute nothing to bending.

E. Material

A single grade is used throughout, defined by a 31-point true stress versus true plastic strain table from a first yield of 360 MPa to 530 MPa at about 5.3% plastic strain, with E = 210 GPa. The identical table was entered in the Abaqus definition. The provenance of this table is itself part of the account: an early, incorrect transcription (a different 385 MPa table) propagated silently through several diagnostic runs before being caught by direct comparison with the source model, an episode analysed in Section IX.

IV. THE HUMAN-AI CO-WORKING METHOD

The workflow follows [1]: a living skill document encoding formulation, benchmark targets, conventions, and known failure modes; formulation-first prompting, in which equations are cited by number rather than described in prose; incremental scope; and mandatory benchmark verification after every stage. This section documents what was specific to the present campaign: a strict division of labour, a pre-registration discipline that made every hypothesis test falsifiable, and the correction loop that the discrepancy of Section V actually followed, including a material-transcription error caught only by the author's direct inspection of the source model.

A. Division of labour

Responsibility split along a consistent line. The LLM implemented every algorithm (the J2 return mapping, both section models, the CR-TL kinematics), generated every Abaqus and shell input file, computed a quantitative prediction of each file's outcome before it was run, and, on receiving results, diagnosed discrepancies and proposed the next hypothesis. The author ran every Abaqus and shell job without exception, extracted every output, verified the material definition directly against the source model's own keyword editor rather than accepting it as given, and, at each verification gate, accepted, rejected, or redirected the LLM's conclusion.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Neither role was delegable to the other: the LLM cannot execute a commercial FEA license or inspect a colleague's original model file, and the author's time is better spent judging results than encoding a Newton radial-return algorithm by hand. Figure 5 summarises the pattern.

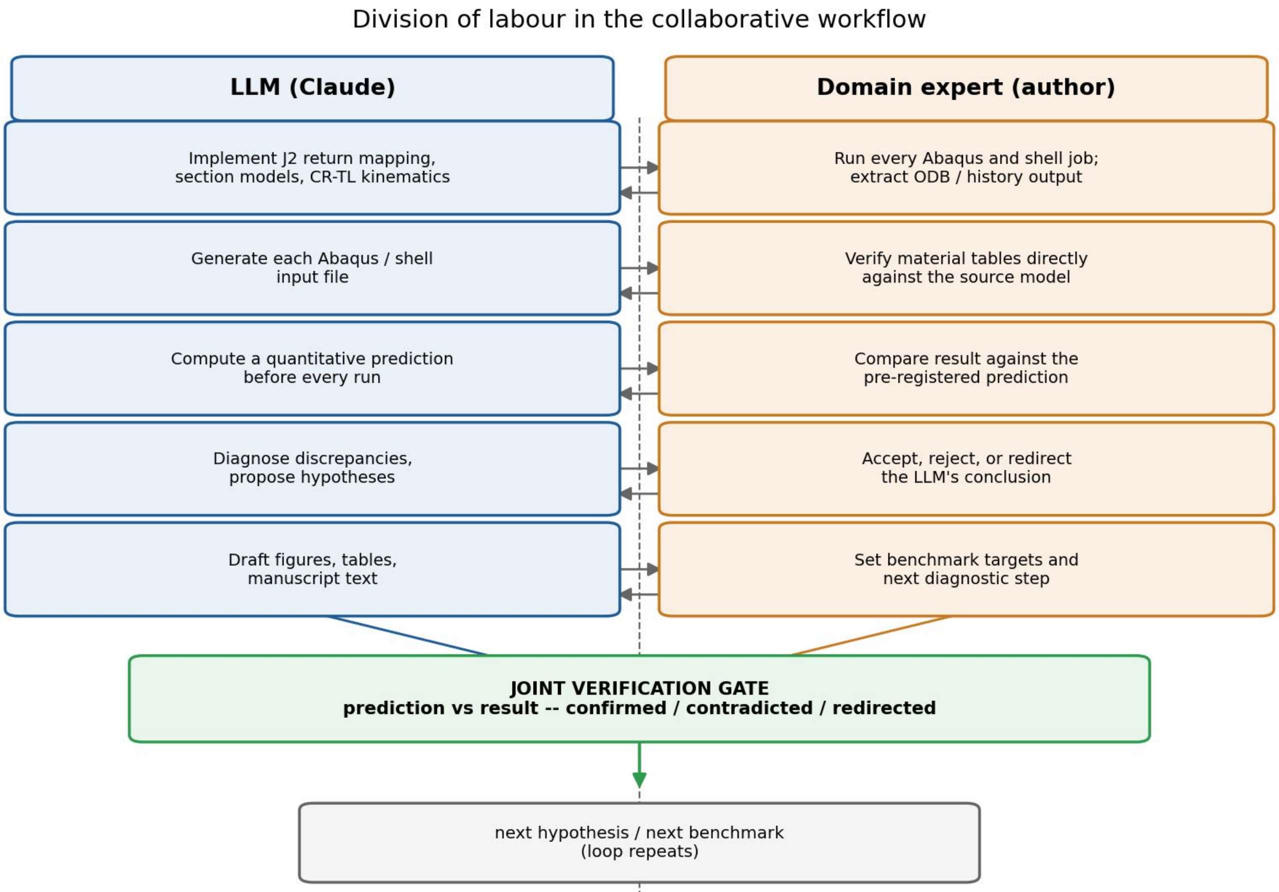

Fig. 3. Division of labour: the LLM implements, predicts, and diagnoses; the author executes every commercial-code run, verifies inputs against source data, and gates each conclusion; results are compared against the pre-registered prediction before the next step is taken.

B. *Pre-registration and the falsifiability discipline*

Every discriminating Abaqus or shell run was preceded by a written, quantitative LLM prediction of its outcome, recorded in the input file's header before the author executed it. This converted each run from an open-ended comparison into a falsifiable test: a result could confirm, contradict, or partially contradict a specific numerical claim, not merely feed a vague impression. Two consequences follow directly in Section VI. A prediction that increasing the PIPE section's circumferential point count would change the tip rotation by a specific amount, made under the (subsequently found incorrect) 385 MPa material, was contradicted by the corresponding run and explicitly retracted rather than quietly revised. A later prediction, that the section-level moment deficit would asymptote toward a specific limit as the section progressively plasticized, was confirmed to better than 0.1% once the material was corrected. Reporting the retraction alongside the confirmation is deliberate: a methodology that only reports its successes cannot be distinguished from one that reports nothing testable at all.

C. *The correction loop*

The investigation did not proceed in a straight line, and Figure 6 records the path actually taken rather than an idealised version of it. Seven of the eight candidate hypotheses of Table 3 were eliminated by ordinary discriminating tests, each following the predict-run-compare cycle of Section IV.B. The eighth candidate, a mismatched material table, was not caught by any discriminating Abaqus run at all: every diagnostic file the LLM had generated used a table transcribed earlier in the project, and every comparison against it was internally consistent, which is precisely why the error was invisible to hypothesis testing.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

It surfaced only when the author, independently and outside the predict-run-compare loop, opened the original source model's keyword editor and compared its * PLASTIC block directly against the table in use. This is the clearest illustration in the paper of why the division of labour in Section IV.A is not a matter of convenience: the LLM had no channel by which to notice that its own working assumption disagreed with a file it had never been shown.

Once corrected, the point-count test was repeated and now showed genuine, monotonic convergence, exposing a further asymmetry: the earlier, uncorrected material had coincidentally suppressed the very effect the test was designed to detect, a false negative that the correction loop caught only because the material was fixed for independent reasons and the test happened to be rerun. The introduction of the shell reference (Section VI.C) was not part of any original plan; it was directed by the author specifically because a quadrature-convergence argument, however clean, could not by itself establish which section idealisation was physically correct, only that a family of them agreed with one another. Anchoring the comparison to a model with no beam-section idealisation at all was the step that converted a self-consistent but potentially circular argument into a validated conclusion.

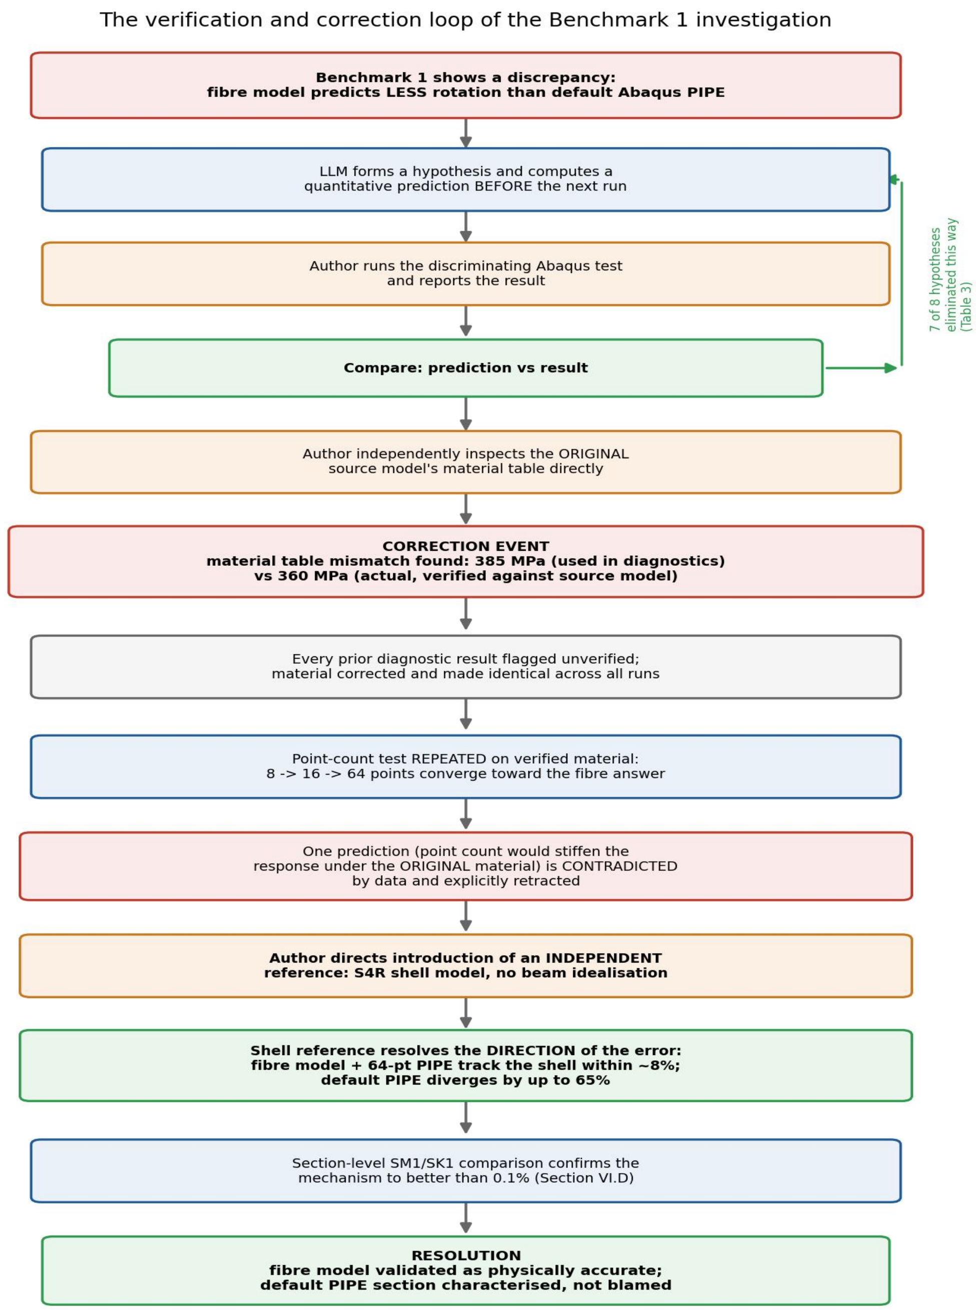

Fig. 4. The verification and correction loop actually followed in the Benchmark 1 investigation

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

D. What this adds to the workflow of [1]

The companion paper characterised omission, syntax, and strategy errors, and identified mandatory benchmark verification as the safeguard against them. The present campaign shows that benchmark verification is necessary but not sufficient on its own: a self-consistent error in a shared input (here, the material table) can pass every internal comparison while still being wrong, and is caught only by a verification channel external to the predict-run-compare loop itself, in this case the author's direct inspection of source data. The extended error taxonomy of Section IX.C generalises this and the other lessons of the correction loop beyond the present case study.

V. BENCHMARK 1 — UNIFORM-MOMENT CANTILEVER

A. Definition

A 10 m cantilever of 16-inch by 21 mm (OD 406.4 mm) pipe carries a concentrated end moment, loaded monotonically through 1300, 1350, 1400, 1450 kNm. First yield of the section is near 800 kNm and the fully-plastic moment near 1600 kNm, so the four levels lie progressively onto the plastic plateau. Under pure end moment the bending moment is constant along the beam, so curvature is uniform and tip rotation equals curvature times length, verified to six decimal places across mesh densities from 2 to 100 elements. A load/unload variant cycles each level back to zero for the path-dependence check.

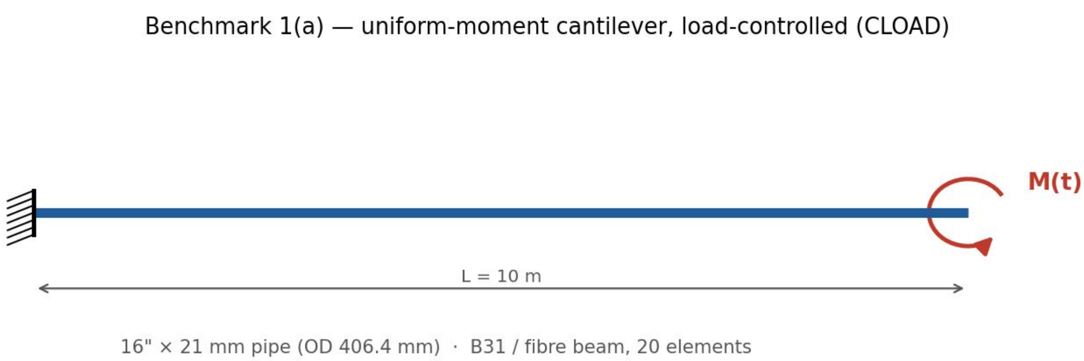

Fig. 3. Load-controlled cantilever, the primary Benchmark 1 model.

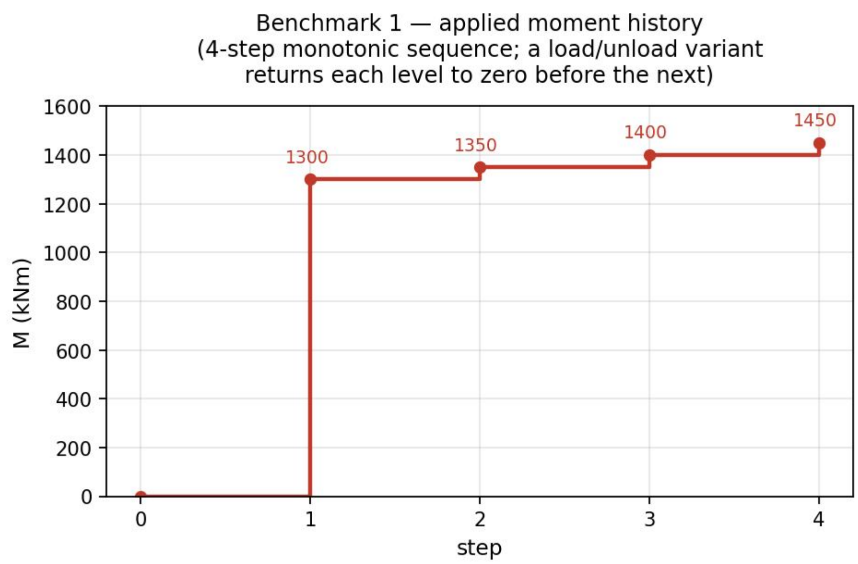

Fig. 4. Applied moment history for Benchmark 1: four-step monotonic sequence, with a load/unload variant returning each level to zero before the next.

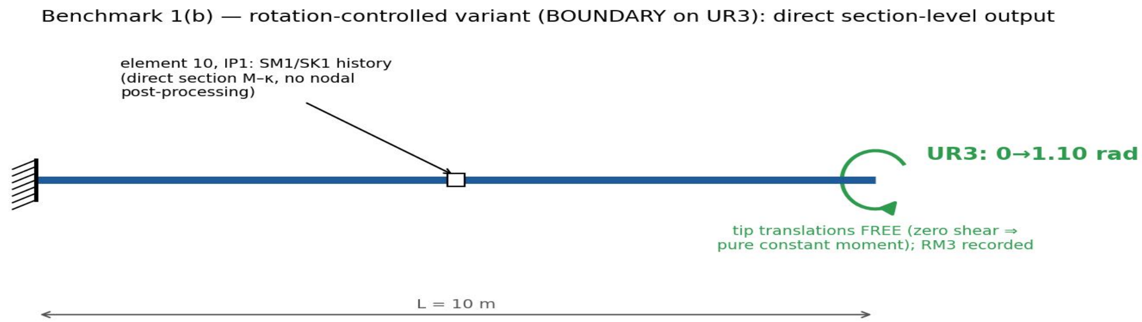

Fig. 5. Rotation-controlled variant of Benchmark 1: tip rotation prescribed directly, translations free, giving zero shear and a pure constant moment; SM1/SK1 recorded at element 10 for the direct section-level comparison of Section VI.D.

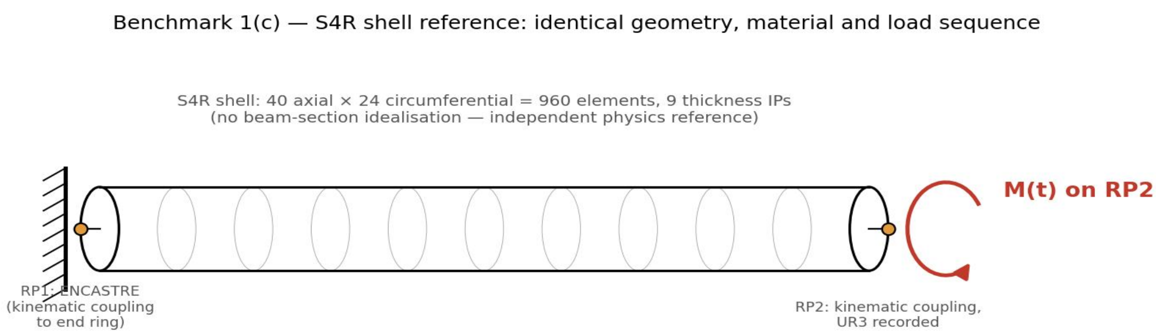

Fig. 6. S4R shell reference for Benchmark 1: identical geometry, material and load sequence, with no beam-section idealisation.

B. Section models and the commercial reference

Table 1 compares tip rotation from the fibre and polar section models against several Abaqus PIPE configurations, all on the identical material. The fibre model predicts markedly less rotation than the default Abaqus PIPE section, the gap growing from 7% at 1300 kNm to 67% at 1450 kNm. Taken alone this looks like an implementation error in the new library. Section VI shows, using an independent shell reference, that the opposite is true.

Table 1. Tip rotation (rad), uniform-moment cantilever, all on the verified material. The default Abaqus PIPE section coincides with the explicit 8-point polar model; refining to 64 points brings Abaqus close to the fibre model; the shell reference sits nearest the fibre / 64-point family, far from the default PIPE section.

M (kNm)	fibre (exact)	polar(8)	Abq default PIPE	Abq 64-pt	shell (S4R)
1300	0.248	0.293	0.293	0.250	0.260
1350	0.319	0.428	0.428	0.321	0.341
1400	0.431	0.667	0.667	0.435	0.467
1450	0.617	1.084	1.080	0.630	0.656

C. Path dependence

Under load/unload the J2 model freezes plastic strain to machine precision at every return to zero moment, with the residual rotation retained; the R-O model returns identically to the undeformed state. This qualitative distinction, confirmed against Abaqus in Section VII for the passage test, is the boundary of validity between the two constitutive paths.

VI. THE INVESTIGATION AND ITS RESOLUTION

A. Framing

The question posed by Table 1 was whether the fibre model was wrong, or the default beam section. The investigation proceeded by eliminating candidate causes and, decisively, by introducing a reference free of any beam idealisation.

B. Material verification and section convergence

The first substantive finding was corrective rather than physical. Several diagnostic runs had inadvertently used an incorrect material table; once the table was verified against the source model and made identical across all runs (Section III.E, Section IX), the fibre-model self-convergence was established (N = 200 strips within 0.01% of N = 1000) and the point-count behaviour of the Abaqus PIPE section could be tested cleanly. Increasing the circumferential point count of the PIPE section from 8 to 16 to 64 produces monotonic convergence of tip rotation toward the fibre-model value (Table 1): at 1450 kNm, 8 points give 1.08 rad, 16 points 0.68 rad, 64 points 0.63 rad, against the fibre model's 0.62 rad. This is the classical signature of quadrature convergence: the default 8-point PIPE section under-integrates the plastic section response, and resolving the circumference removes the error.

This point-count sensitivity did not appear under the earlier incorrect material, whose particular curve shape masked it, an object lesson recorded in Section IX.

C. The shell reference resolves the direction of the error

A quadrature convergence argument shows the section models are mutually consistent but does not, by itself, establish which is physically correct. An S4R shell model of the same cantilever, carrying no beam-section idealisation, settles it. Table 2 gives the error of each beam result against the shell reference. The fibre model and the 64-point PIPE section track the shell to within 4-8% across the whole load range, a roughly constant offset consistent with ordinary beam-versus-shell theory (the beam cannot capture cross-section ovalisation under large plastic curvature). The default PIPE section, by contrast, departs from the shell reference increasingly with load, reaching about 65% at 1450 kNm.

Table 2. Tip-rotation error against the independent S4R shell reference. The fibre model and the well-resolved 64-point section stay within a stable single-digit band; the default 8-point PIPE section diverges increasingly with load.

M (kNm)	fibre vs shell	64-pt vs shell	default PIPE vs shell
1300	-4.5%	-3.8%	+12.7%
1350	-6.6%	-5.9%	+25.5%
1400	-7.7%	-6.9%	+42.8%
1450	-5.9%	-4.0%	+64.6%

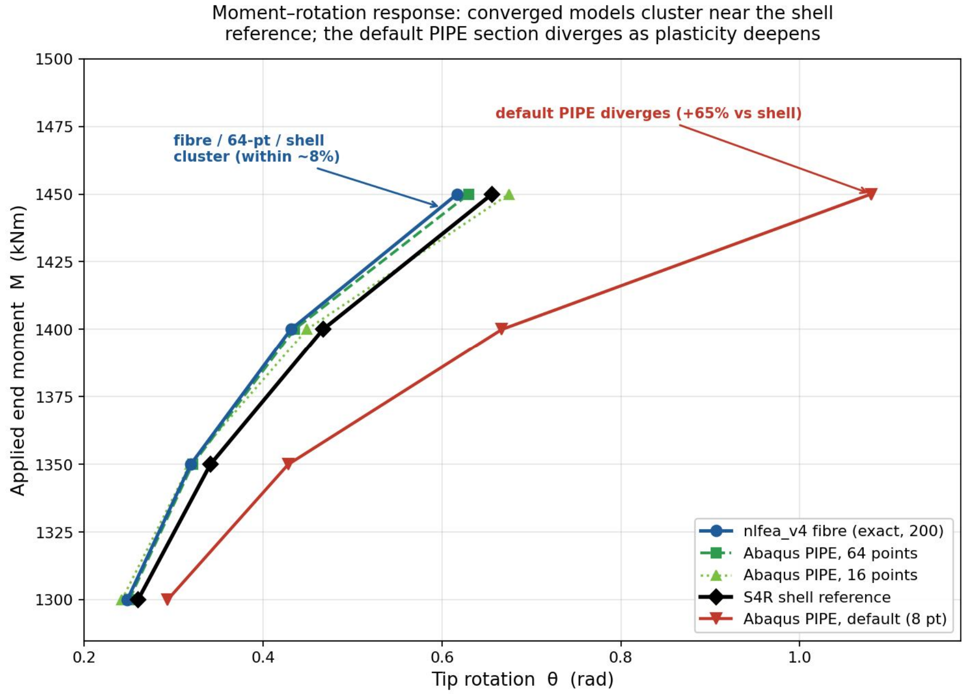

Fig. 7. Moment-rotation response – Bench Mark Test 1

D. Section-level confirmation

A rotation-controlled run (tip rotation prescribed, translations free so shear is zero) provided the section moment-curvature relation directly from Abaqus section output, bypassing all nodal post-processing. At matched curvature the fibre-model moment exceeds the default-section moment by only 2.4% at low curvature growing to 4.7% at high curvature, the small constitutive-level difference that the shallow post-yield moment-curvature slope amplifies into the large tip-rotation gap of Table 1. Comparing at matched curvature rather than matched moment is therefore essential, a metric point developed in Section IX.

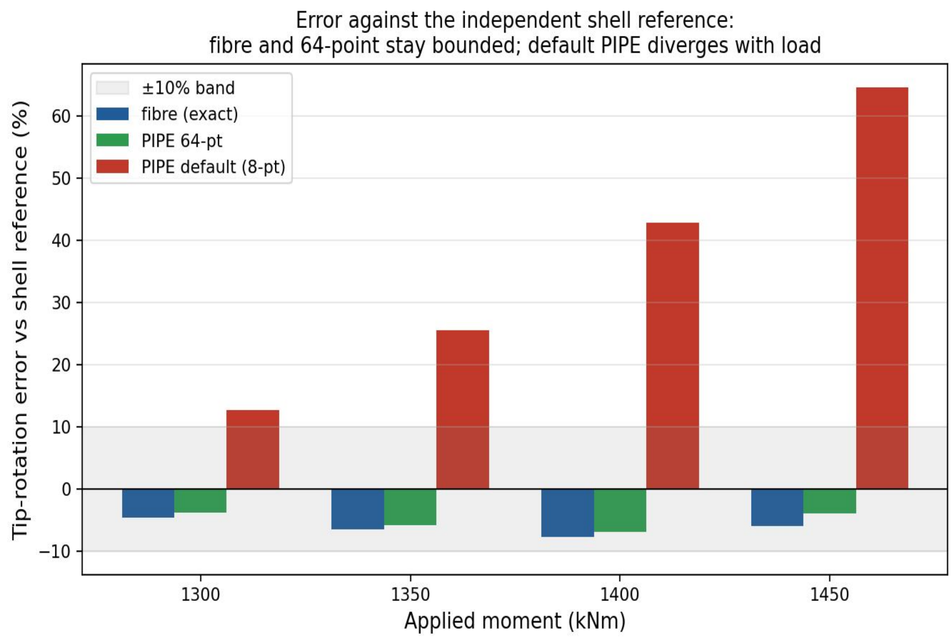

Fig. 8. Error against the shell reference

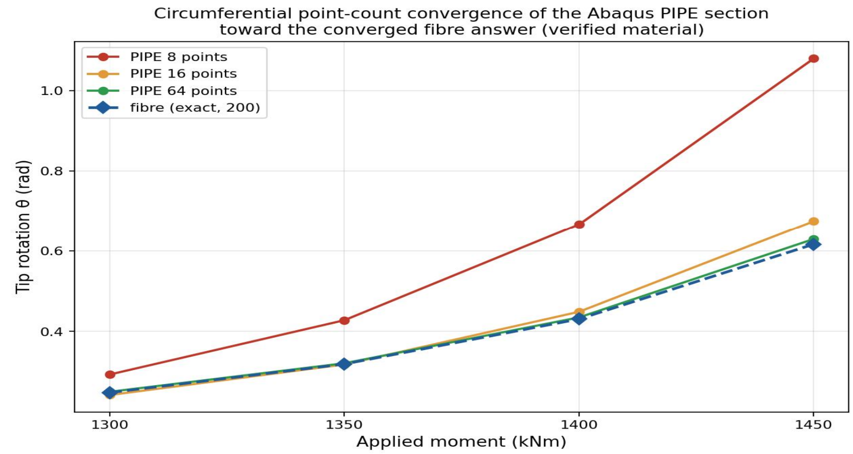

Fig. 9. Circumferential point-count convergence of the Abaqus PIPE section toward the fibre answer

E. The interchangeable compatibility section

For users who must reproduce an existing default-PIPE Abaqus model, the polar section with 8 aligned points reproduces the default Abaqus PIPE tip rotation to within about 1% across all four load levels (Table 1). It is retained in the library as an Abaqus-compatibility mode; the fibre model remains the accurate default, now shell-validated.

F. Open item

One observation remains unexplained and is reported as such. The default (no options) PIPE run and the original benchmark model, though both nominally default and now on identical material, do not coincide (about 0.29 versus 0.27 rad at 1300 kNm). The difference is not section point count, which is identical; it is most likely a structural difference (mesh density, load application, or boundary-condition scheme) between the simplified diagnostic template and the original full model, and is flagged for future resolution. It does not affect the shell-validated conclusions above.

VII. BENCHMARKS 2 AND 3 — MODERATE PLASTICITY

A. Five-point displacement-controlled bend

A 32 m pipe is bent over five stations at 8 m intervals, the three interior stations displacement-driven onto a circular arc progressively tightened to R = 100, 85, 70, 60 m. Two sections (16-inch and 20-inch) were run against matching Abaqus models. Peak strain occurs at the quarter-span stations, not midspan, at 1.4-2.1 times the nominal strain, confirming the discrete-support concentration relevant to stinger design. At these strain levels (nominal 0.19-0.40%) the section works on the steep part of the moment-curvature curve and no amplification operates. Table 3 shows both section models agreeing with Abaqus in low single digits; the polar section tracks nominal strain to within about 0.5% while the fibre model runs a stable 4-5% low, and both track plastic strain within a few percent (excepting a near-yield point where the small denominator inflates the percentage).

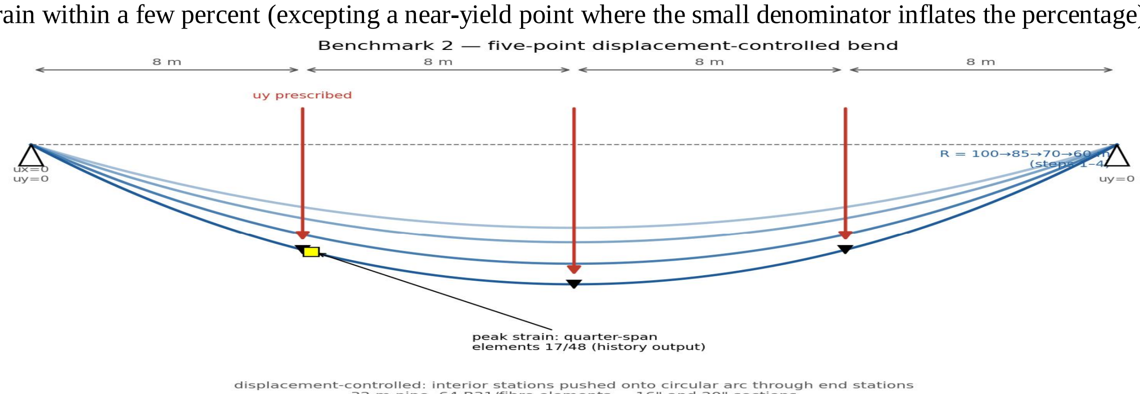

Fig. 8. Five-point displacement-controlled bend model: interior stations pushed onto progressively tightening circular arcs (R = 100 to 60 m); peak strain occurs at the quarter-span elements, not midspan.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Table 3. Five-point bend: peak nominal (NE) and plastic (PE) strain error against Abaqus at the quarter-span element, verified material. (*) At 16in/R=100, PE is barely above first yield; the percentage is unstable and the absolute difference is negligible (Section IX.A).

Section	R (m)	NE err exact / polar	PE err exact / polar
16in	100	+0.1% / +0.5%	(near-yield*)
16in	85	-4.1% / +0.1%	+4.8% / +4.8%
16in	70	-4.7% / -0.1%	+3.8% / +4.2%
16in	60	-5.3% / -0.5%	+2.9% / +3.4%
20in	100	-4.1% / +0.4%	+4.1% / +5.2%
20in	85	-4.5% / -0.0%	+3.7% / +4.0%
20in	70	-5.1% / -0.4%	+2.8% / +3.2%
20in	60	-4.9% / -0.5%	+3.2% / +3.0%

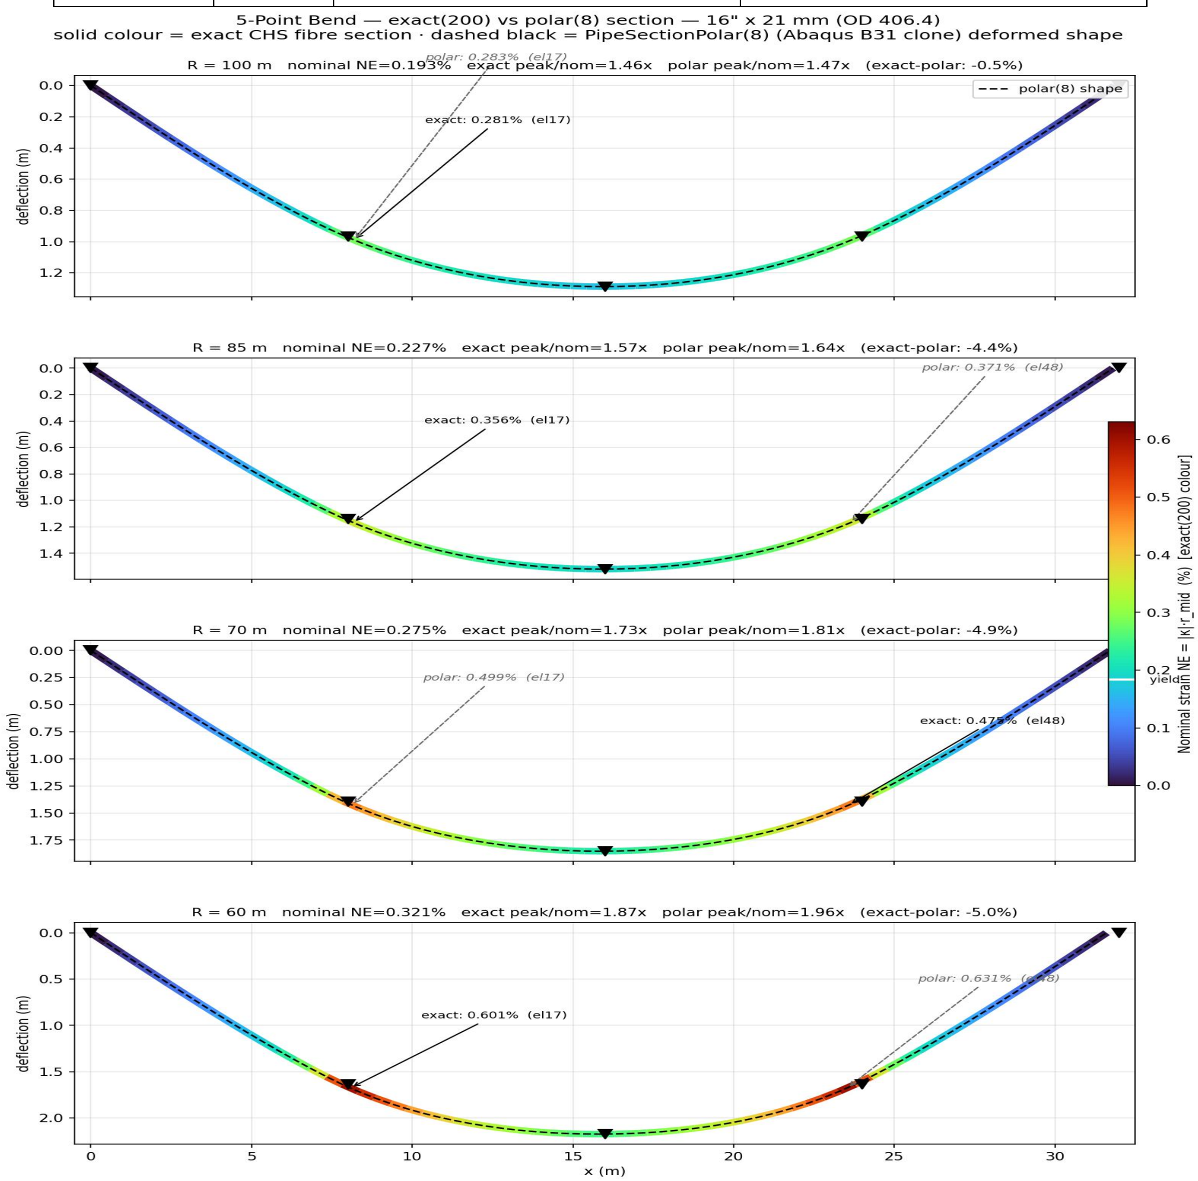

Fig. 9. Five-point bend: deformed shapes coloured by nominal strain, four radii, 16-inch. Peak strain occurs at the quarter-span stations; the polar-section shape is overlaid dashed.

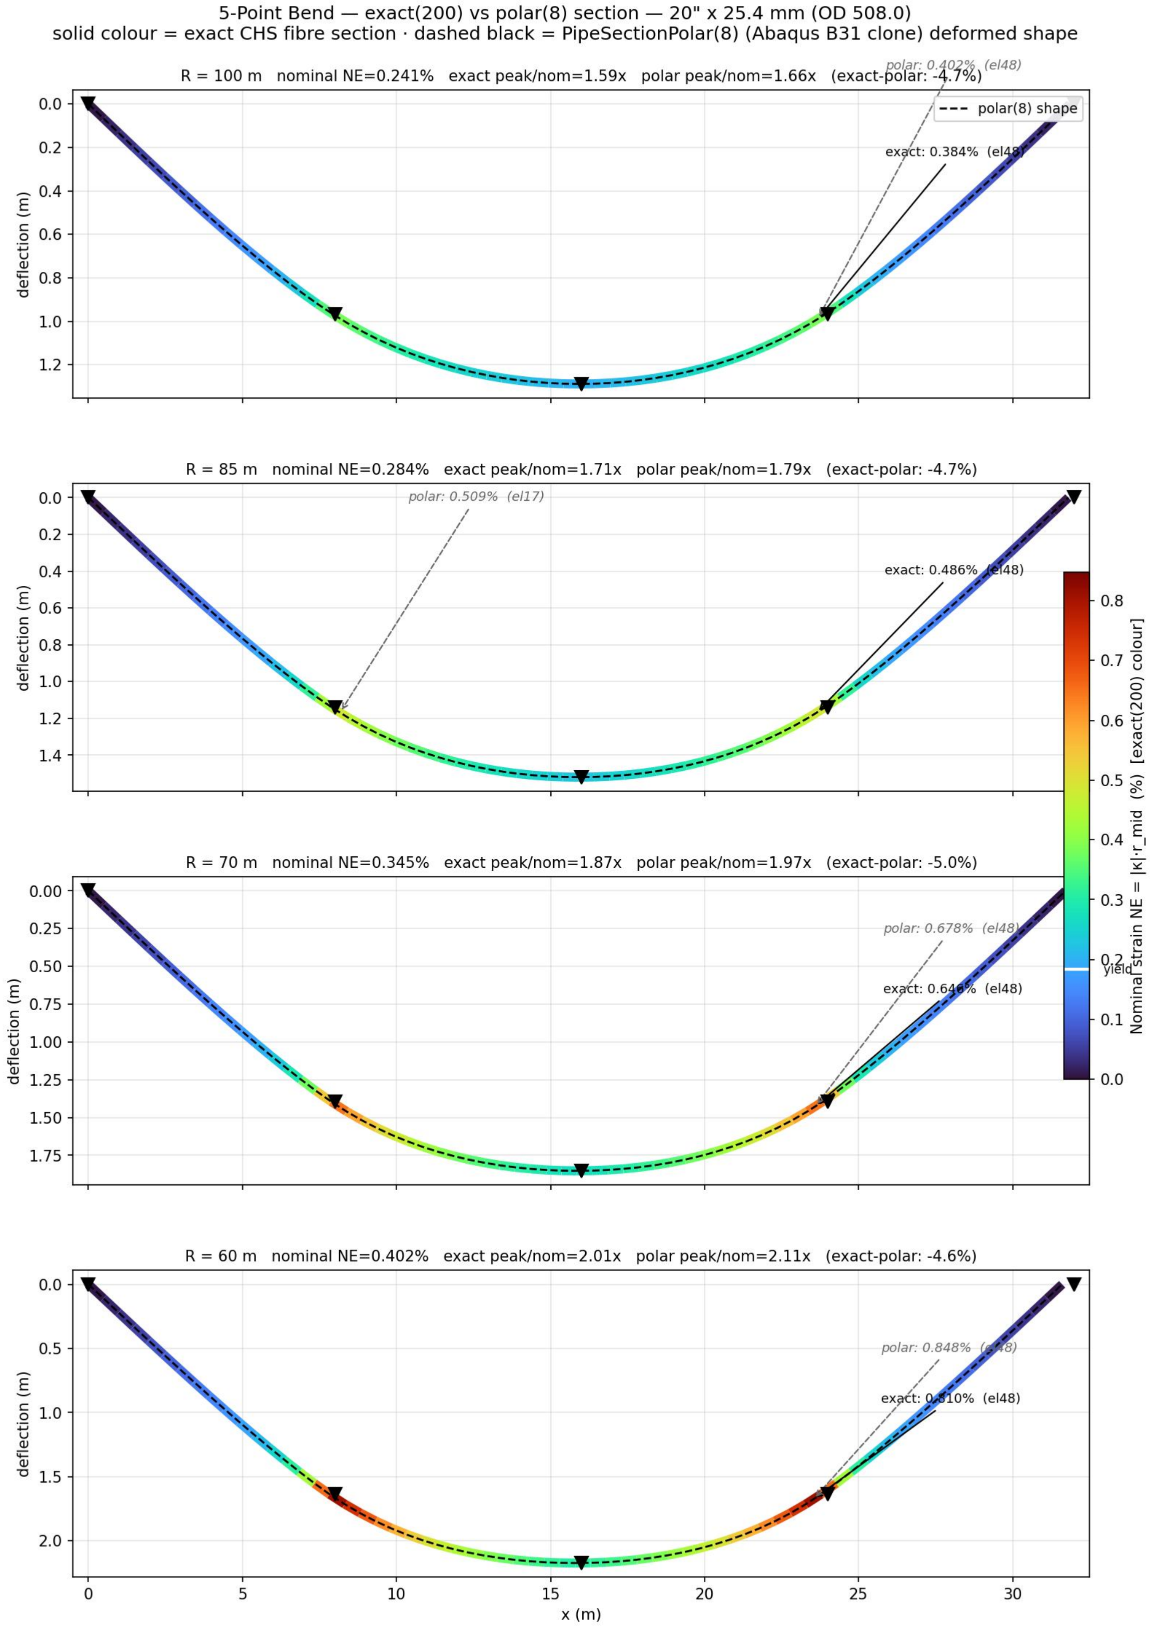

Fig. 9. Five-point bend: deformed shapes coloured by nominal strain, four radii, 20-inch (lower). Peak strain occurs at the quarter-span stations; the polar-section shape is overlaid dashed.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

B. Moving-support passage test

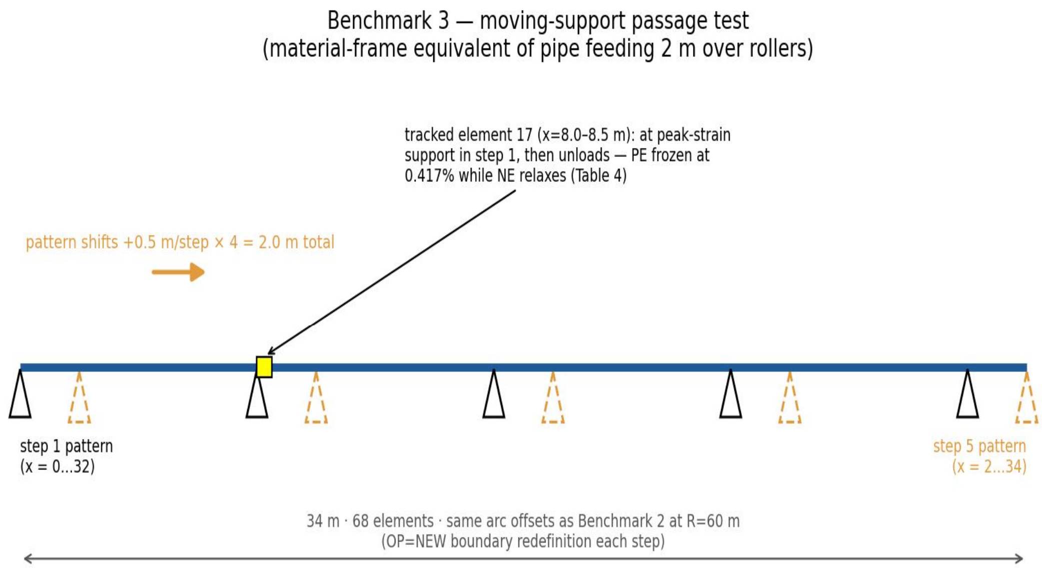

Fig. 10. Moving-support passage model: the five-station pattern translated 0.5 m per step, the material-frame equivalent of a pipe feeding 2 m over stinger rollers; the tracked element sits at the peak-strain support in step 1.

The five-station pattern (R = 60 m) is translated along a stationary 34 m pipe by one element per step, four steps, the material-frame equivalent of a pipe feeding 2 m over rollers. The tracked material point, at the peak-strain quarter support in step 1, unloads as the pattern departs: its nominal strain relaxes about 12% over the four steps while its plastic strain remains frozen. This freezing is confirmed independently in Abaqus, whose plastic strain at the tracked element holds constant at 0.417% across all five steps, matching the fibre model (0.429%, frozen) and the polar model (0.431%, frozen). Nominal strain agreement follows the same pattern as the five-point bend: the polar model within 0.5-2.3%, the fibre model a stable 5-8% low. A deformation-plasticity model run through the same history discards the residual state entirely, the quantity that accumulated-strain assessment requires.

Table 4. Passage test, tracked element 17. Plastic strain is frozen to the reported precision in all three models across every step; nominal strain relaxes as the support pattern departs.

Step	NE_abq	NE exact / polar err	PE_abq	PE exact / polar (frozen)
1	0.00634	-5.3% / -0.5%	0.00417	0.00429 / 0.00431
2	0.00618	-6.9% / -2.2%	0.00417	0.00429 / 0.00431
3	0.00598	-7.2% / -2.3%	0.00417	0.00429 / 0.00431
4	0.00582	-7.5% / -2.3%	0.00417	0.00429 / 0.00431
5	0.00569	-7.7% / -2.3%	0.00417	0.00429 / 0.00431

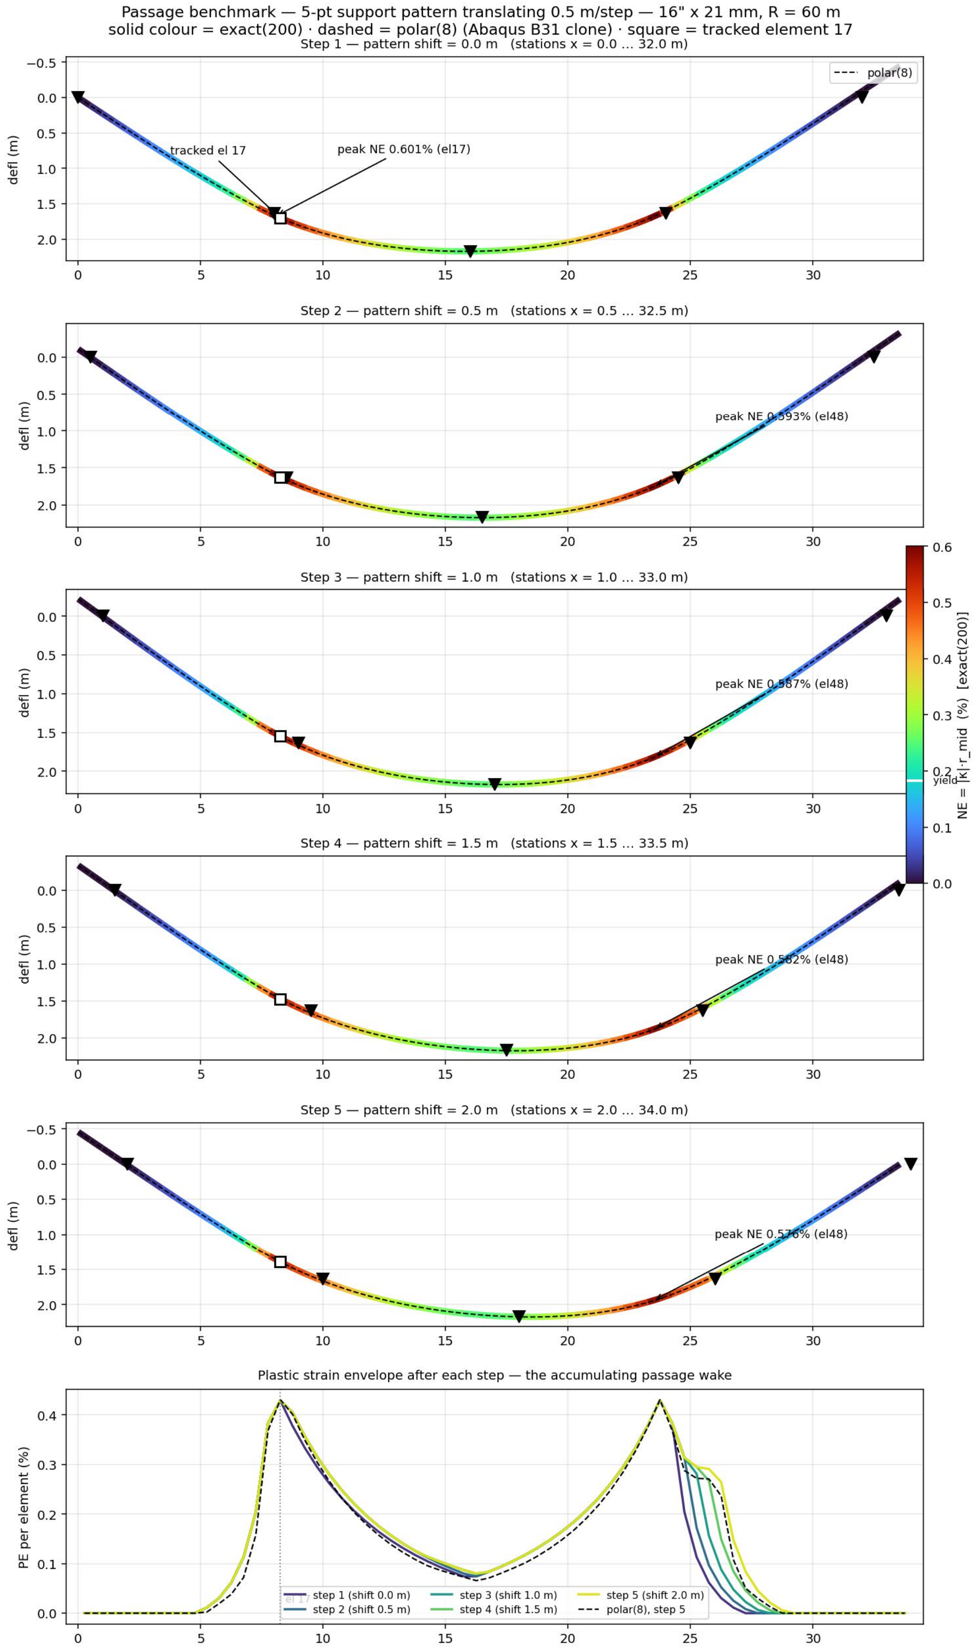

Fig. 11. Passage test: the five-station pattern translating 0.5 m per step (panels 1-5), and the plastic-strain envelope after each step forming the accumulating wake (bottom panel). The tracked element is marked with a square.

VIII. DEFORMATION VERSUS INCREMENTAL PLASTICITY — VALIDITY BOUNDARY

Running both constitutive paths through Benchmark 1 fixes the boundary concretely. On monotonic loading the two agree closely (within about 4% at every level, the R-O path marginally softer as its smooth curve yields earlier than the multilinear table). Under load/unload they separate qualitatively: the J2 model retains the residual rotation and frozen plastic strain confirmed against Abaqus, while the R-O model returns identically to zero. The R-O path is therefore appropriate where DNV practice applies it, single-pass monotonic strain screening; the J2 path is required wherever material points experience load reversal, including roller passage and multi-stage installation, subject to the isotropic-hardening limitation of Section X.

IX. DISCUSSION

A. Two metric and definition pitfalls

Two lessons generalise beyond this study. First, material definitions must be verified against the actual source model, not assumed: an incorrect table (385 MPa first yield instead of 360 MPa) propagated silently through several diagnostic runs, shifting results by 15% or more, and was invisible until the source model's keyword editor was inspected directly. Worse, its particular curve shape coincidentally suppressed the point-count sensitivity of Section VI.B, briefly leading the investigation to a false negative. Second, in deep plasticity, compare codes on moment-curvature at matched curvature, not on rotation or deflection at matched load; near the plastic plateau the inverse reading amplifies constitutive-level differences by the reciprocal of the shallow curve slope, here turning 2-5% into 14-65%, and can misdirect an entire investigation. Near first yield, compare plastic strain in absolute terms, since it is the small difference of two nearly equal quantities and an absolute agreement of 0.0003 can masquerade as a large percentage error.

B. Implications for practitioners

For pipe-bending plasticity benchmarked against Abaqus B31/PIPE, the default section is elastically exact and accurate through moderate plasticity (the entire Benchmark 2 regime), but under-integrates the plastic bending response increasingly as the section approaches its fully-plastic moment, diverging from a shell reference by up to about 65% in rotation near M/Mp of 0.9. Where deep-plateau response matters, the circumferential integration should be refined (64 points effectively converges) or a shell reference used. The compatibility section published here reproduces the default behaviour within about 1% for those matching an existing model.

C. Extended error taxonomy

Four error classes observed here extend the taxonomy of [1]:

1) Formulation-simplification errors: the one-step radial return, correct for the linear-hardening special case silently assumed, wrong for the tabular case actually required. Detectable only by a tabular-curve benchmark.
2) Material-definition errors: a mis-transcribed hardening table, silent at runtime, propagating through many derived files, and detectable only by direct comparison with the source model. The most consequential error of the campaign, and a caution that data provenance deserves the same verification discipline as code.
3) Verification-tooling errors: a curvature-extraction formula in a diagnostic script (not the solver) produced a spurious result that briefly indicted a correct solver; verification code needs the same scrutiny as production code.
4) Hypothesis-management errors: a section-quadrature hypothesis was prematurely conceded on one run that had not, in fact, tested what it appeared to, and was reinstated only when section-level data and the shell reference clarified matters. Pre-registered predictions made both the false step and the correction legible.

D. Division of labour

The pattern of [1] repeats with sharper edges. The LLM implemented every algorithm, generated every input file, and performed every diagnostic derivation; the author ran the commercial and shell references, extracted their outputs, and, more than once decisively, refused to accept a premature conclusion and directed the introduction of the shell arbiter. One exchange is worth recording: when the author questioned whether per-element rotation errors accumulate over a multi-element beam, the disagreement was resolved not by assertion but by an algebraic argument (relative errors of summed proportional quantities do not compound) and a mesh-independence demonstration (identical tip rotation from 2 to 100 elements). Productive disagreement of this kind is a feature of the workflow.

X. LIMITATIONS AND FUTURE WORK

1) Isotropic hardening only; the Bauschinger effect is not represented, so deep strain-reversal problems (reeling, repeated reverse bends) require kinematic or combined hardening. The present load/unload benchmarks do not enter reverse yield.
2) Two-dimensional, bending-dominated; combined axial tension with bending, the true S-Lay overbend state, is the named next step.
3) The passage travel of 2 m demonstrates frozen plastic strain under unloading but not a fully released residual wake; a longer-travel run is a trivial extension.
4) The unresolved default-versus-original-model discrepancy of Section VI.F, a structural rather than sectional difference, remains open.
5) Single-expert, single-model case study, as in [1].

XI. CONCLUSION

1) The open beam FEA library of [1] now carries verified material plasticity: a Ramberg-Osgood path for monotonic screening and a J2 incremental path, with Newton radial return, fibre section integration, and committed plastic state, for path-dependent analysis. Plastic strain freezes exactly on unloading, confirmed against the commercial code.
2) Against an independent shell reference, the converged fibre-integration section tracks physical reality to within 4-8% across the full load range, while the default Abaqus B31/PIPE section under-integrates the plastic bending response, diverging by up to about 65% in rotation near the fully-plastic moment. A point-count study (8, 16, 64) confirms this as a quadrature-convergence effect.
3) Across a multi-support bend matrix and a moving-support passage test at moderate plasticity, both section models agree with Abaqus in low single digits, and the compatibility section reproduces the default Abaqus PIPE response within about 1%.
4) Two definition and metric pitfalls, silent material-table mismatches and matched-load rather than matched-curvature comparison, are documented as the dominant hazards of the campaign, alongside an extended four-class error taxonomy for AI-assisted numerical software development.
5) Domain-expert-guided AI collaboration, with mandatory benchmark verification, pre-registered predictions, and, critically, an independent physics reference, produced a verified plasticity extension and corrected a natural misreading of the commercial code, not by trusting either implementation but by anchoring both to shell-level physics.

XII. ACKNOWLEDGEMENTS AND DATA AVAILABILITY

The author thanks Anthropic for development of the Claude large language model. No external funding was received; the author declares no conflicts of interest.

The library (including both section models), all benchmark and plotting scripts, all Abaqus and shell input files, and the figures of this paper will be archived at the repository given in [1]. All numerical results are reproducible from the repository.

REFERENCES

[1] Sivaraman S M 2026 AI-assisted development of a nonlinear finite element analysis program: a case study using a large language model Int. J. Res. Appl. Sci. Eng. Technol. 14 253-264, DOI 10.22214/ijraset.2026.83397
[2] Sivaraman S M 2015 A VBA based computer program for nonlinear FEA of large displacement 2D beam structures Int. J. Res. Appl. Sci. Technol. 3 1-21
[3] Ramberg W and Osgood W R 1943 Description of stress-strain curves by three parameters NACA Technical Note 902
[4] DNV 2021 DNV-ST-F101 Submarine pipeline systems (Hovik: DNV)
[5] Simo J C and Hughes T J R 1998 Computational Inelasticity (New York: Springer)
[6] Crisfield M A 1991 Non-linear Finite Element Analysis of Solids and Structures vol 1 (Chichester: Wiley)
[7] Bathe K J 1996 Finite Element Procedures (Englewood Cliffs: Prentice Hall)
[8] De Souza R M 2000 Force-based finite element for large displacement inelastic analysis of frames PhD thesis University of California Berkeley
[9] Holsgrove S C and Lyons L P R 1989 Benchmark Tests for Two-Dimensional Thin Beams and Axisymmetric Shells NAFEMS Report N4
[10] Bai Y and Bai Q 2014 Subsea Pipeline Design, Analysis and Installation (Amsterdam: Gulf Professional Publishing)
[11] Dassault Systemes 2024 Abaqus Analysis User's Guide: Beam cross-section library (Providence: Dassault Systemes Simulia)
[12] Hou X et al 2024 Large language models for software engineering: a systematic literature review ACM Trans. Softw. Eng. Methodol. (in press) arXiv:2308.10620

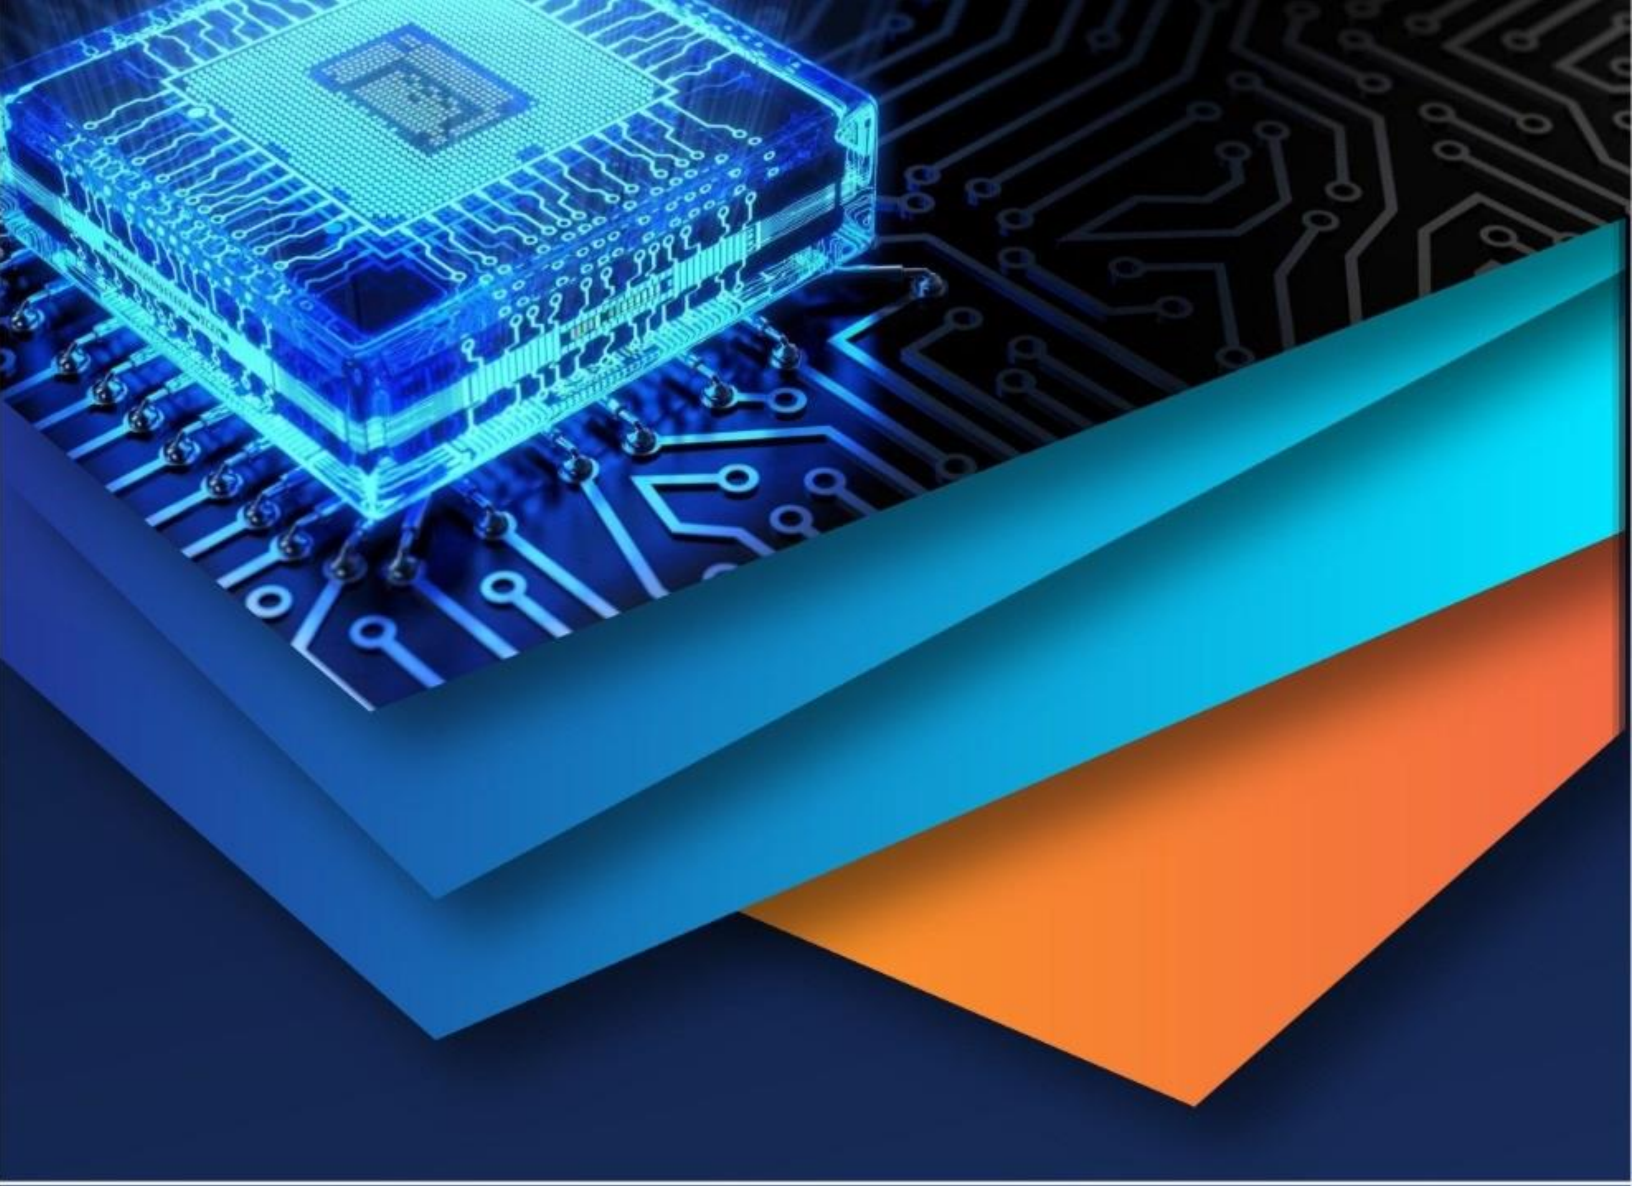

10.22214/IJRASET

INDEX COPERNICUS
45.98

ISRA JIF
IMPACT FACTOR:
7.129

TOGETHER WE REACH THE GOAL
IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY