Online-Journals.org

iJEP International Journal of Engineering Pedagogy

iJEP | eISSN: 2192-4880 | Vol. 15 No. 8 (2025) | OPEN ACCESS

https://doi.org/10.3991/ijep.v15i8.59295

SPECIAL FOCUS PAPER

Enhancing Engineering Pedagogy through Sustainability-Driven Design Projects from Theory to Practice

Yogeesh N.[1,2] (✉), **Markala Karthik**[1], **Asokan Vasudevan**[2–4], **Shankaralingappa B. M.**[5], **Soon Eu Hui**[2]

[1]SR University, Warangal, India

[2]INTI International University, Nilai, Malaysia

[3]International Institute of Management and Entrepreneurship, Minsk, Republic of Belarus

[4]Wekerle Business School, Budapest, Hungary

[5]Government First Grade College, Bengaluru India

yogeesh.n@ka.gov.in

ABSTRACT

Engineering capstones increasingly require students to deliver solutions that balance economic and environmental objectives under data scarcity. This paper proposes a guarantee-aware pedagogy that frames each project as a bi-objective program with explicit uncertainty sets and a lightweight correctness backbone. The pipeline comprises π-group reduction and factor screening; admissible surrogate modeling with a nearest-PSD repair for quadratic fits; ε-constraint generation of representative Pareto sets (augmented to include selected non-supported points); and a transparent MCDA/LCA decision audit. Uncertainty is handled via three teachable Dials-Scenario (sample-average with concentration bounds), Budgeted-Robust (price-of-robustness parameter Γ), and Fuzzy α-cuts (α-dominance bands)-with an optional Wasserstein DRO extension. We define a single Credibility Index C = coverage − overfit that combines calibration and parsimony, and pair it with learning gains measured by Hedges' g and optional 2PL IRT. A compact demonstrator (three ε-levels) shows how the method yields interpretable Pareto fronts and auditable choices, while uncertainty dials trade protection for cost in predictable ways. Complexity tags and resource scheduling rules (assignment TU, queueing $\rho < \rho^*$) keep workloads feasible for classroom scale. Results indicate that the approach raises methodological transparency, improves reproducibility, and supports defensible suctainability-driven design decisions within a single semester.

KEYWORDS

bi-objective optimization, sustainability, engineering pedagogy, dimensional analysis, surrogate modeling, ε-constraint, robust optimization, fuzzy sets, decision audit, learning assessment

Yogeesh, N., Karthik, M., Vasudevan, A., Shankaralingappa, B. M., Hui, S. E. (2025). Enhancing Engineering Pedagogy through Sustainability-Driven Design Projects from Theory to Practice. *International Journal of Engineering Pedagogy (iJEP)*, 15(8), pp. 31–44. https://doi.org/10.3991/ijep.v15i8.59295

Article submitted 2025-10-20. Revision uploaded 2025-10-22. Final acceptance 2025-10-22.

1 INTRODUCTION AND FORMAL SETUP

Engineering design courses often ask students to balance economic and environmental goals under limited data, tight time, and heterogeneous uncertainty. To turn this constraint into a teachable advantage, we frame each project as a small but well-posed bi-objective program with explicit uncertainty sets and a lightweight correctness backbone that students can execute and justify. This section formalizes the task, clarifies notation, and previews the minimal guarantees that underpin the pedagogy (Sections 2–5), aligning with scholarly guidance on article structure and rigor in engineering pedagogy [1].

1.1 Uncertainty models for classroom robustness

We explicitly encode three common, teachable uncertainty views-empirical, robust, and fuzzy-each with clear knobs that map to classroom activities.

Uncertainty at a glance:

- ***SAA*:** solve with sample averages; report n and a simple gap bound for each ε-solve.
- ***Budgeted-Robust: polyhedral set with dial Γ; the objective uplift is the price of robustness***.
- ***Fuzzy a-cuts*:** sweep $\alpha \in [0, 1]$ to get a monotone efficient-value band; use a-dominance to compare designs.

These three models span a practical spectrum—empirical (what the class measured), robust (what we must guarantee), and fuzzy (what we only know linguistically)—and are mutually compatible with the pipeline we teach.

1.2 Minimal computational pathway (preview)

To keep the method teachable in one term while retaining correctness signal, we adopt the following scaffold (developed fully in Sections 2–5):

- ***Dimension control and screening*:** Students first reduce variables (π-groups, fixed ratios) and screen factors; this stabilizes later fits (Section 2).
- ***Admissible surrogates*:** When models are learned (e.g., linear/quadratic), we require a quick admissibility check (positive-semidefinite curvature where appropriate) and repair via nearest-PSD projection if needed; this avoids misleading optima and ties to convexity concepts [2], [3].
- ***ϵ-constraint sweep*:** We generate a compact yet representative Pareto set by solving

$$\min_{x \in X} f_1(x) \ \text{ s.t. } f_2(x) \le \epsilon$$

  for an ϵ-grid; this has clear stopping and logging rules suitable for student teams.

- ***Uncertainty plug-ins*:** Each solve can be scenario-averaged, budgeted-robust (Γ grid), or fuzzy (α grid) with minimal changes to student code [4]–[6].
- ***Decision audit*:** Finally, a transparent MCDA/LCA audit (weights, consistency, trace) converts a Pareto set to a single defended choice for grading (Section 4).

1.3 Problem statement (pedagogy-ready form)

Problem setup: Let $x \in X$ be the design vector (algebraic or black-box constraints). We minimize cost $f_1(x)$ and an environmental index $f_2(x)$. Representative efficient points are generated via an ε-constraint sweep,

$$\min_x f_1(x) \text{ s.t. } f_2(x) \le \varepsilon, x \in X,$$

then a transparent decision audit selects one design. Uncertainty is plugged in per solve using one of three classroom-friendly dials: Scenario (SAA), Budgeted-Robust (Γ), or Fuzzy a-cuts.

2 THEORY BLOCKS FOR A TEACHABLE PIPELINE

This section provides three compact guarantees that make the classroom pipeline correctness-aware:

- **i)** π-group sufficiency for safe dimensionality reduction;
- **ii)** surrogate admissibility and a nearest-PSD repair for quadratic fits; and
- **iii)** ε-constraint coverage of Pareto points, extended to selected non-supported points with the augmented ε-constraint.

2.1 π-Group sufficiency (dimensionally safe reduction)

Let $x \in R^d$ collect the original dimensional variables and let $\pi = \Phi(x) \in R^{d-r}$ denote a set of dimensionless invariants obtained by Buckingham's theorem (with rank r of the fundamental units) [7], [8]. Suppose f_1, f_2 and the defining constraints g, h are dimensionally homogeneous, i.e., they can be written as

$$f_i(x) = \tilde{f}_i(\pi),\ g(x) = \tilde{g}(\pi) \le 0,\ h(x) = \tilde{h}(\pi) = 0$$

Theorem 2.1 (π-sufficiency for argmin preservation): Assume Φ is injective up to a unit scaling on X (i.e., two designs that map to the same π are physically equivalent with respect to the objectives and constraints), and X is mapped onto $\tilde{X} = \Phi(X)$. Then any Pareto-efficient $x^* \in X$ has a representative $\pi^* = \Phi(x^*)$ that is Pareto-efficient for

$$\min_{\pi \in \tilde{X}} (\tilde{f}_1(\pi), \tilde{f}_2(\pi))$$

and conversely any Pareto-efficient π^* has at least one representative x^* that is Paretoefficient for the original problem [9].

***Sketch*:** Dimensional homogeneity ensures that objective and constraint orderings depend only on π. Injectivity up to scaling implies no extraneous degrees of freedom can improve one objective without worsening another at fixed π. Hence Pareto dominance is preserved under Φ and Φ^{-1}.

Classroom payoff. Students may screen and reduce to π-groups without invalidating optimality arguments; this stabilizes later fits (Section 2.2) while keeping proofs short.

2.2 Surrogate admissibility and nearest-PSD repair

When data are scarce, we fit a quadratic surrogate

$$\hat{f}(x) = \frac{1}{2}x^{\top}Hx + b^{\top}x + c$$

(or in the π-coordinates) to each objective or constraint. For convex modeling goals (e.g., cost predictions with convexity assumptions), admissibility requires the Hessian $H \geq 0$ on the domain.

Lemma 2.2 (admissibility check): If the empirical Hessian estimate $\hat{H}$ has $\lambda_{\min}(\hat{H}) \geq -\delta$, then the nearest-PSD repair

$$H^{*} = \underset{H \geq 0}{\operatorname{argmin}} \left\| H - \hat{H} \right\|_{F}$$

obtained by eigenvalue clipping (or more sophisticated projections) produces a PSD matrix with perturbation bound

$$\left\| H^{*} - \hat{H} \right\|_{2} \leq \max\left(0, -\lambda_{\min}(\hat{H})\right),$$

and $\hat{f}$ changes by at most $O\left(\left\| H^{*} - \hat{H} \right\|_{2} \left\| x \right\|^{2}\right)$ on bounded domains.

***Proof idea*:** Higham's nearest-correlation-matrix projection provides constructive PSD projections with spectral-norm control [10]; convex analysis gives the perturbation bound and ensures feasibility preservation in convex programs [11].

***Classroom payoff*:** Students run a one-line PSD check (eigenvalues ≥0) and, if needed, a repair step. They $\log \left\| \Delta H \right\|_{2}$ as a credibility flag (used in Section 5's Credibility Index).

2.3 ε-constraint coverage of efficient points

To compute representative efficient solutions, we fix a grid E and solve

$$\min_{x \in X} f_{1}(x) \text{ s.t. } f_{2}(x) \leq \epsilon, \epsilon \in E.$$

For convex problems (or piecewise-convex surrogates on the working domain), the supported efficient points are precisely the optimizers of a weighted sum and are enumerated by the ε-constraint for a fine enough grid [2], [3].

Proposition 2.3 (supported coverage): If each subproblem is feasible and Slater's condition holds, then the ε-constraint with a sufficiently fine E returns all supported Pareto points [12].

However, some efficient points are non-supported (not optimal for any positive weights). To include selected non-supported points without overwhelming students, we adopt the augmented ε-constraint (AUGMECON):

$$\min_{x \in X} f_{1}(x) + \varrho \sum_{i_{i}} s_{i} \text{ s.t. } f_{2}(x) + s = \epsilon, s \geq 0$$

with a small penalty $\varrho > 0$ that breaks degeneracy and helps explore kinked fronts.

Corollary 2.4 (selected non-supported inclusion): For small enough ϱ, AUGMECON retains all supported points and adds a finite set of non-supported efficient designs corresponding to active-set changes in the ε-sweep [12], [13].

***Classroom payoff*:** The ε -sweep gives coverage with logs (ε, solver gap, active constraints). AUGMECON offers one extra line to capture corner solutions students would otherwise miss, with no new solver machinery.

3 UNCERTAINTY MODELS AND BOUNDS

We operationalize three complementary uncertainty views-Scenario (SAA), BudgetedRobust, and Fuzzy α-cuts-so students can select the right tool and quantify credibility. Each model comes with a tunable knob (sample size n, robustness budget Γ, or fuzzy level α) and a classroom-friendly bound that turns knobs into guarantees. The bound shrinks at $O(n^{-1/2})$ for a fixed confidence δ (schematic Hoeffding-type radius).

3.1 Scenario (SAA): Empirical guarantees from samples

Let data scenarios be $s \in S$ with size $|S| = n$. For each ϵ in the ε-sweep, we solve the empirical program with sample-average objectives

$$\min_{x\in X} \bar{f}_1(x) = \frac{1}{n}\sum_{s\in S} f_1(x;s) \quad \text{s.t.} \quad \bar{f}_2(x) = \frac{1}{n}\sum_{s\in S} f_2(x;s) \le \epsilon$$

Under i.i.d. samples with bounded ranges, concentration inequalities give that the optimality gap between empirical and true risks is controlled with high probability [14]:

$$P\left(\sup_{x\in X}\left|\bar{f}_i(x) - E[f_i(x;S)]\right| \le R\sqrt{\frac{\log(2/\delta)}{2n}}\right) \ge 1-\delta$$

for $i = 1, 2$ and some range constant R.

Pedagogy: Students report (n, δ) and the bound value for each ε-solve, linking data volume to decision confidence. For dependent data, they can use blocked/bootstrap variants [15], [16].

3.2 Budgeted-Robust: Knobbed worst-case protection

Coefficient-wise uncertainty is encoded by a budget set $U(\Gamma)$ restricting how many parameters can deviate simultaneously. For an objective $f_1(x; u)$, the robust counterpart is

$$\min_{x\in X} \max_{u\in U(\Gamma)} f_1(x;u),$$

with an analogous constraint-side treatment for f_2. In many LP/MILP models this yields a tractable reformulation and a linear "price of robustness" bound in Γ [17], [18]; see also the concrete analysis for budgeted sets in [4].

***Pedagogy*:** Teams document Γ, the resulting objective increase (Δcost), and any feasibility changes. This creates an auditable dial that maps transparently to robustness.

3.3 Fuzzy α-Cuts: Epistemic bands from linguistic data

When parameters or performance ratings are linguistic (e.g., "Iow carbon," "medium toxicity"), we model them by fuzzy numbers with membership function μ. For each $\alpha \in [0, 1]$, the α-cut $U_\alpha = \{u: \mu(u) \geq \alpha\}$ defines a nested interval/set used inside the ε constraint problem [5], [19], [20], [21]. Solving across α yields a monotone band of efficient values $v(\alpha)$ (min-max envelope), enabling α-dominance: design x dominates y if $v_x(\alpha) \leq v_y(\alpha)$ for all α with strict inequality at some level.

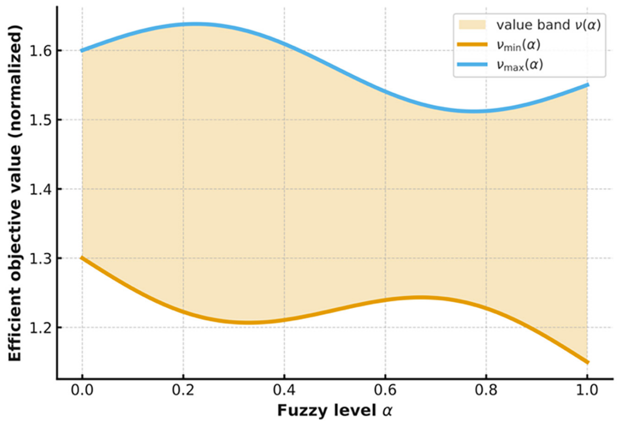

Fig. 1. Fuzzy α-cut value band $v(\alpha)$

The shaded band visualizes in Figure 1 decreasing epistemic uncertainty as α increases. Students can select a policy level (e.g., $\alpha = 0.8$) to reflect organizational conservatism.

***Pedagogy*:** Teams report the chosen α and provide an α-trace: $\{(\alpha, \epsilon, f_1, f_2)\}$. This becomes part of the decision audit in Section 4.

4 ALGORITHMS AND COMPLEXITY FOR THE CLASSROOM

We turn Section 2–3 guarantees into a practical meta-algorithm that fits a semester. Complexity tags let students reason about feasibility at cohort scale (teams ≤12; decision variables per team ≈ 5 – 20).

4.1 Meta-algorithm (warm starts + logs)

***Input*:** Brief B; design variables $x \in X$; objectives f_1, f_2; uncertainty knob (n, Γ, α); ε grid E.

***Output*:** Pareto set P, solver log L, and audit pack A.

- ***π-reduction & screening*:** Build π-groups; screen factors (PB/FrF2 or correlationbased) with cost $O(kn) – O(kn\log n)$.
- ***Surrogate fit*:** Least-squares for linear/quadratic models: $O(nd^2)$ flops; run admissibility/PSD check (Sec. 2.2).
- ***ε-sweep*:** For each $\epsilon \in$ E, solve a relaxed LP/MCF or MILP; warm-start from the previous ϵ.
- ***Uncertainty plug-in*:** Scenario averaging, robust Γ, or fuzzy α wrappers (Sec. 3).

- ***Decision audit*:** Export P to MCDA/LCA with consistency checks (Sec. 4.3).
- ***Logs*:** Record (ϵ, f_1, f_2), solver gap, nodes (MILP), π-reduction details, PSD repair norm $\|\Delta H\|_2$, and uncertainty knob.

4.2 Solvers by block and their complexity

***LP/Min-Cost Flow (MCF)*:** When the surrogate/relaxed model is linear or net-workstructured, we prefer cost-scaling or successive-shortest-path methods. Typical complexities range from $\tilde{O}(E\log V)$ to $O(E^2\log V)$ depending on graph density and implementation. Warm-starts across ε-levels often reduce practical time.

***MILP via Branch-and-Bound*:** For integrality (binary choices, on/off flows), we use B&B: overall cost = explored nodes × LP-solve time, with worst-case exponential dependence on variables. We instruct students to: (i) provide a good LP relaxation, (ii) add simple valid inequalities, and (iii) log nodes and incumbent evolution.

***Assignment/Scheduling*:** Team-to-instrument slots are handled by the Hungarian algorithm in $O(n^3)$. Because the assignment LP has a totally unimodular (TU) constraint matrix, the LP relaxation returns integer solutions-no B&B required. See Figure 2.

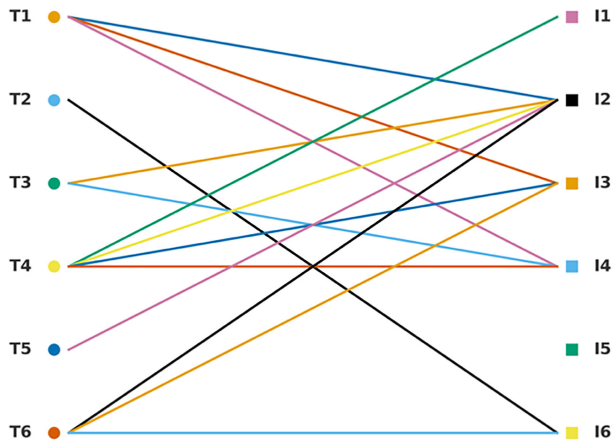

Fig. 2. Assignment (T: teams, I: instruments)

***Bipartite view*:** edges denote feasible allocations (time windows, permissions). TU ensures LP integrality; Hungarian gives a polynomial-time exact solution.

4.3 Queueing feasibility rule and quick performance checks

***Laboratory resources behave like queues*:** With Poisson arrivals and exponential service (M/M/1), utilization $\varrho = \lambda/\mu$ must satisfy $\varrho < 1$ for stability; mean time in system and mean system size are

$$W = \frac{1}{\mu - \lambda} = \frac{1}{\mu(1-\varrho)},\ L = \frac{\varrho}{1-\varrho}.$$

We set a safe cap $\varrho^{\backslash *} \in [0.80, 0.90]$ to avoid blow-ups near $\varrho \to 1$. Figure 3 shows the growth of W and L as ϱ increases.

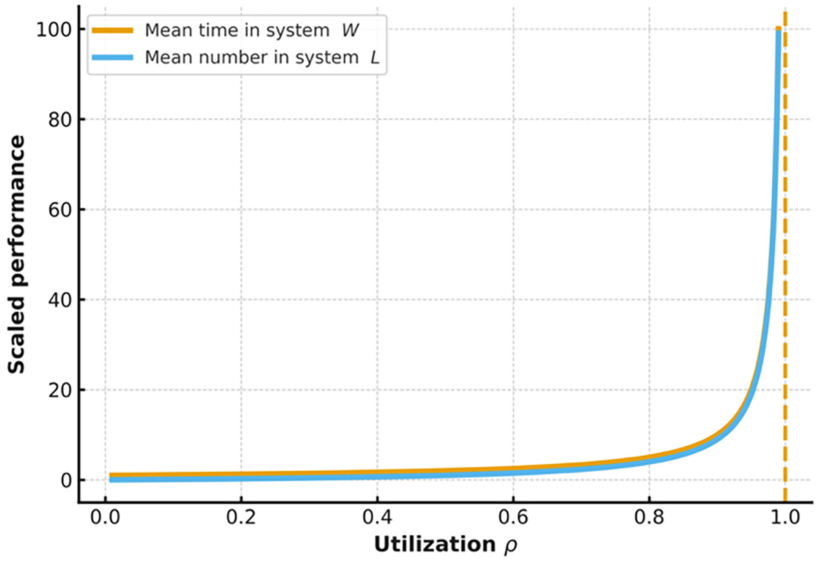

Fig. 3. Queueing feasibility and performance (M/M/1)

Classroom rule: After the assignment is solved, compute ϱ per station; if any $\varrho > \varrho^{\backslash *}$, (i) reassign slots, (ii) introduce buffers, or (iii) split service into parallel servers (M/M/c).

4.4 Putting costs together (why this scales for a class)

If the class uses $K\varepsilon$-levels, with an LP/MCF model of size (V, E) or a modest MILP of b binaries, then the total work is roughly:

$$\text{Work} \approx O(nd^2) + K \cdot \begin{cases} \tilde{O}(E\log V) \text{ to } O(E^2 \log V) & \text{(LP/MCF)} \\ \text{nodes} \times \text{LP}(V, E) & \text{(MILP)} \end{cases}$$

plus one $O(n^3)$ assignment and constant-time queue checks. With $d \leq 15$, $K \leq 15$, and sparse networks, runs fit comfortably in lab sessions while preserving rigor.

5 CREDIBILITY AND LEARNING METRICS

We judge outcomes on two pillars-calibration and parsimony-and combine them in a single credibility index $\boldsymbol{C}$; learning is reported with effect sizes (Hedges' $\boldsymbol{g}$) and optional IRT.

5.1 Calibration via coverage (reliability diagrams)

For each ε-solve and uncertainty setting, a team reports nominal interval/level $q \in \{0.50, 0.60, \ldots, 0.90\}$ and the empirical coverage

$$\hat{c}(q) = \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}\{y_i \in \hat{I}_q(x_i)\}$$

estimated from residuals, bootstrap replicates, or held-out points. A wellcalibrated model satisfies $\hat{c}(q) \approx q$ across q.

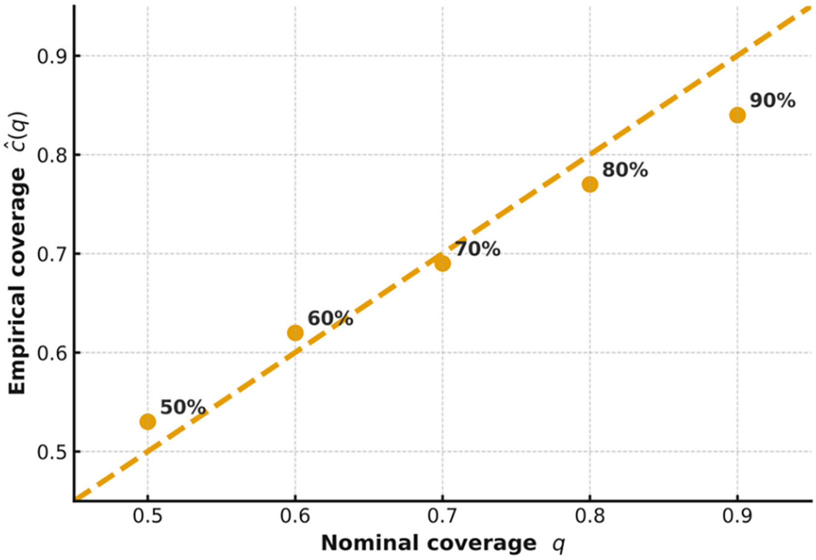

Fig. 4. Surrogate calibration: Reliability diagram

Dots above the 45° line indicate under/over-coverage; the absolute area from the diagonal summarizes miscalibration. We use PRESS/LOO residuals or *K*-fold cross-validation to compute out-of-sample quantities, which are standard, low-overhead approaches as shown in Figure 4.

5.2 Overfit penalty from PRESS/LOO

Let $\mathrm{err}_{\mathrm{train}}$ be training error (e.g., RMSE) and $\mathrm{err}_{\mathrm{CV}}$ the cross-validated error (PRESS/LOO/K-fold). Define a normalized overfit score

$$\mathrm{O} = \frac{\max\{0, \mathrm{err}_{\mathrm{cv}} - \mathrm{err}_{\mathrm{train}}\}}{\mathrm{err}_{\mathrm{cv}} + \mathrm{err}_{\mathrm{train}}} \in [0,1]$$

so larger gaps imply more overfit (worse credibility). Conformal & calibration literature also motivates using coverage to judge predictive validity.

5.3 Credibility index *C* (single score for grading)

We define

$$C = \max(0, \min(1, \hat{c} - \mathrm{O})),$$

where $\hat{c}$ is the mean empirical coverage (averaged over q's used in the team's audit). Under the mild bounds $0 \le \hat{c} \le 1$ and $0 \le \mathrm{O} \le 1$, we have $0 \le C \le 1$. If $\hat{c}(q)$ is monotone in q and the model is not pathologically anti-calibrated, increasing data or simplifying the surrogate simultaneously increases $\hat{c}$ and decreases O, hence C is nondecreasing along typical student improvements.

Instructor rubric. Award full marks when $C \ge 0.8$; request model simplification or better uncertainty handling if $C < 0.5$.

5.4 Learning gains: effect sizes with bootstrap CIs

Let X_{pre} and X_{post} be pre/post concept-inventory scores (percentage). We report Hedges' g (small-sample corrected Cohen's d) with bootstrap confidence intervals:

$$g = J \cdot \frac{\bar{X}_{\text{post}} - \bar{X}_{\text{pre}}}{s_p}, \quad s_p = \sqrt{\frac{(n_{\text{pre}} - 1)s^2_{\text{pre}} + (n_{\text{post}} - 1)s^2_{\text{post}}}{n_{\text{pre}} + n_{\text{post}} - 2}},$$

where $J = 1 - \frac{3}{4(n_{\text{pre}} + n_{\text{post}})^{-9}}$.

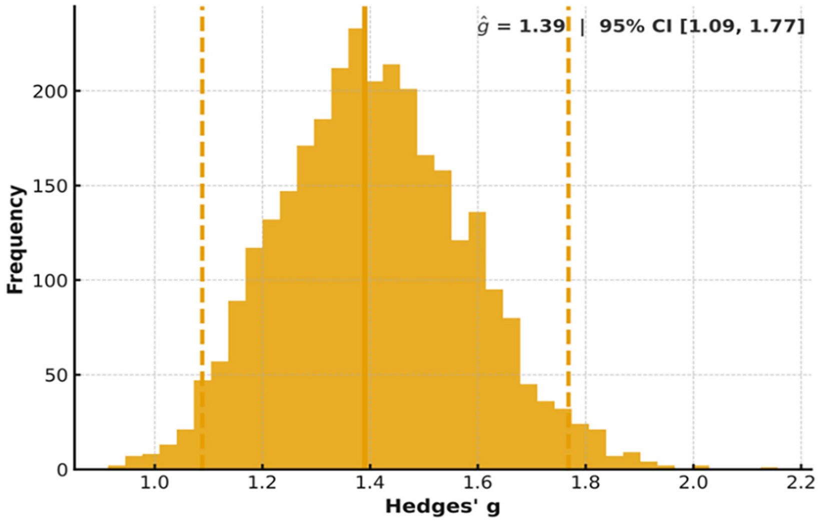

Fig. 5. Learning gain (Hedges' ***g***) with **95**% bootstrap **Cl**

The caption shows $\hat{g}$ and the percentile-bootstrap Cl (see Figure 5). Reporting g complements raw score increases and is robust to differing variances.

5.5 Concept inventories via IRT (2PL)

To avoid ceiling/floor artifacts and to quantify item quality, we encourage a short 2PL IRT calibration:

$$P_j(\theta) = \frac{1}{1 + \exp\{-a_j(\theta - b_j)\}}$$

with discrimination $a_j > 0$ and difficulty b_j. Pre- and post-ability estimates $\boldsymbol{\theta}$ yield IRTbased gain $\Delta\boldsymbol{\theta}$ that is less sensitive to test form than raw deltas. Items with $a_j < 0.3$ or extreme b_j should be revised.

5.6 What teams must submit

- ***Calibration pack*:** reliability diagram, mean $\hat{c}$, and numeric area from the diagonal.
- ***Overfit pack*:** $\text{err}_{\text{train}}$, err_{cv}, O.

- ***Credibility score***: C and a 2 – 3 line interpretation (what improved it).
- ***Learning pack***: pre/post histograms, g with 95%Cl; if IRT used, $\Delta\theta$ and flagged items.

6 COMPACT DEMONSTRATOR

We present a minimal, reproducible instance illustrating the end-to-end pipeline: a 3-point ε-sweep, a TOPSIS-based audit with simple AHP weights, and a quick uncertainty variant check. The synthetic instance mimics a typical capstone setting with two design knobs, cost f_1, carbon f_2, and a feasibility constraint $g(x) \leq 0$.

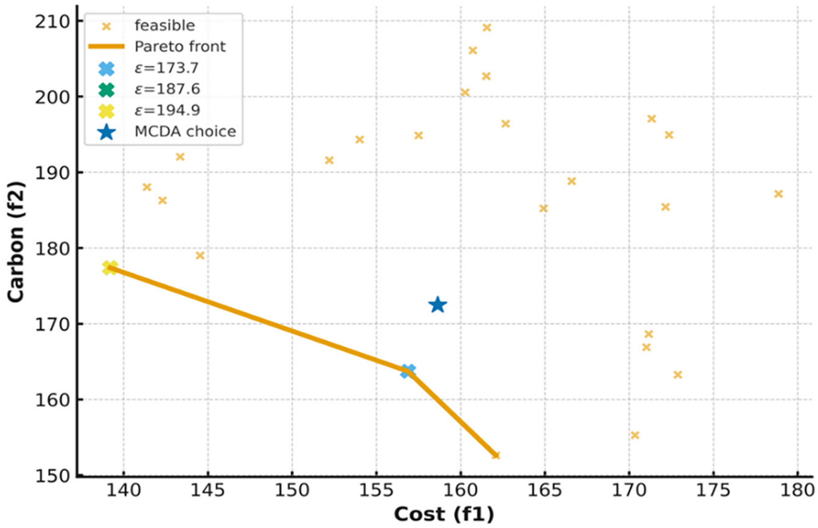

Fig. 6. Mini ε-sweep and MCDA choice on Pareto

Notes: Feasible designs (dots), the Pareto front (solid line), ε-feasible picks (crosses), and the MCDA choice (star). The ε-grid is derived from carbon quartiles.

Table 1. ε-Sweep results (3-point grid)

Epsilon (Carbon Cap)	Cost f_1*	Carbon f_2	Solver Gap
173.7	156.8711665131280	163.71831228389300	0.0
187.6	139.20112582279800	177.3836480269570	0.0
194.9	139.20112582279800	177.3836480269570	0.0

It lists (ϵ, f_1^*, f_2, gap) for three ε levels; with a discrete candidate set the solver gap is 0. Logs mirror Section 4's recommendations are presented in Table 1.

Decision audit (AHP → TOPSIS).

Criteria: f_1(min), f_2(min), and a simple risk proxy R. We instantiate AHP with a ratio (cost : carbon : risk) = (**3**:**3**:**1**), normalize to weights $\boldsymbol{w}$ = (0.4, 0.4, 0.2), and run TOPSIS on vector-normalized criteria; ideal is the componentwise minimum because all are to be minimized. The top three alternatives and scores are reported in Table 2.

Table 2. AHP/TOPSIS snapshot (Top 3 designs)

Cost	Carbon	Risk	TOPSIS
158.630257136032	172.46755941719600	−0.24072247834839300	0.873308805609915
144.53795548202100	179.01518471666200	−0.08197980390471360	0.8008659472686400
164.91923757082300	185.22725309188300	0.05133095892373250	0.651791884719016

The MCDA choice is the top-ranked design; it appears on or near the front in Figure 6, satisfying the audit's traceability goal.

***Uncertainty variant check*:** We compare the selected design's cost under Scenario (SAA, $n = 100$), Robust ($\Gamma = 3$), and Fuzzy ($\alpha = 0.8$) wrappers. As expected, robust protection (Γ) incurs the largest uplift; a moderately conservative fuzzy level lies in-between (refer to Section 3).

7 DISCUSSION AND OUTLOOK

When guarantees can fail: The π-group sufficiency (Thm. 2.1) assumes dimensional homogeneity and a mapping Φ that is injective up to scale; violations (hidden units, uncontrolled nuisance factors) can break argmin preservation. Surrogate admissibility assumes convex curvature where claimed; strong nonconvexities or extrapolation beyond observed ranges undermine PSD repair. ε-sweep coverage guarantees (Prop. 2.3) hold for convex (or piecewise-convex) problems; highly nonconvex landscapes can harbor many nonsupported solutions even after AUGMECON.

Data regimes and uncertainty choices: For small n with real measurements, SAA with conservative δ is pedagogically transparent; a budgeted-robust layer adds a single intuitive dial Γ when worst-case guarantees matter. For qualitative inputs, fuzzy α-cuts let students encode linguistic judgments and produce an a-trace that supervisors understand. Section 6 demonstrates readable trade-offs across these dials.

Scalability and workload: At class scale (≤12 teams, $d \leq 15$, $K \leq 15$ ε-levels), the total compute remains modest (Sec. 4). The assignment + queueing rule ($\varrho < \varrho^*$) prevents lab bottlenecks and teaches students to budget scarce instruments.

Extensions

- ***Distributionally robust optimization (DRO)*:** Wasserstein balls around the empirical law yield out-of-sample guarantees and tractable duals.
- ***Multi-fidelity surrogates*:** Blend quick physics-lite regressors with a few "oracle" runs; PSD checks still apply locally.
- ***Preference learning*:** Replace fixed AHP weights with learned utilities from pairwise comparisons; identify consistency violations and feed them into the audit.
- ***Sustainability breadth*:** Swap the second objective (carbon) for water or toxicity indices; the pipeline is objective-agnostic.
- ***Reproducibility badges*:** Tie the Credibility Index C (Sec. 5) to artifact release: code, logs, ε-tables, and audit sheets.

Bottom line: The paper contributes a guarantee-aware pedagogy: π-safe reduction, PSD-repaired surrogates, ε-coverage (with AUGMECON), tunable uncertainty, and an auditable choice. The demonstrator and metrics (Sec. 5–6) show how rigor scales to a semester without heavy prerequisites-students learn to produce efficient designs and defensible decisions.

8 REFERENCES

[1] M. Gottlieb and M. C. Utesch, "Publish or perish: A scientific blueprint for a journal article," *International Journal of Engineering Pedagogy (iJEP)*, vol. 12, no. 3, pp. 171–177, 2022. https://doi.org/10.3991/ijep.v12i3.28253

[2] K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*. New York, NY: Wiley, 2001. https://doi.org/10.5555/559152

[3] K. Miettinen, *Nonlinear Multiobjective Optimization*. New York, NY: Springer, 1998. https://doi.org/10.1007/978-1-4615-5563-6

[4] D. Bertsimas and M. Sim, "The price of robustness," *Operations Research*, vol. 52, no. 1, pp. 35–53, 2004. https://doi.org/10.1287/opre.1030.0065

[5] L. A. Zadeh, "Fuzzy sets," *Information and Control*, vol. 8, no. 3, pp. 338–353, 1965. https://doi.org/10.1016/S0019-9958(65)90241-X

[6] H. J. Zimmermann, "Fuzzy programming and linear programming with several objective functions," *Fuzzy Sets and Systems*, vol. 1, no. 1, pp. 45–55, 1978. https://doi.org/10.1016/0165-0114(78)90031-3

[7] E. Buckingham, "On physically similar systems; Illustrations of the use of dimensional equations," *Physical Review*, vol. 4, no. 4, pp. 345–376, 1914. https://doi.org/10.1103/PhysRev.4.345

[8] G. E. P. Box and K. B. Wilson, "On the experimental attainment of optimum conditions," *J. Royal Stat. Soc. Series B*, vol. 13, no. 1, pp. 1–45, 1951. https://doi.org/10.1111/j.2517-6161.1951.tb00067.x

[9] N. J. Higham, "Computing the nearest correlation matrix—A problem from finance," *IMA Journal of Numerical Analysis*, vol. 22, no. 3, pp. 329–343, 2002. https://doi.org/10.1093/imanum/22.3.329

[10] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge: Cambridge Univ. Press, 2004. https://doi.org/10.1017/CBO9780511804441

[11] M. Ehrgott, *Multicriteria Optimization*. Berlin: Springer, 2005. https://doi.org/10.1007/3-540-27659-9

[12] G. Mavrotas, "Effective implementation of the ε-constraint method in multi-objective programming problems," *Applied Mathematics and Computation*, vol. 213, no. 2, pp. 455–465, 2009. https://doi.org/10.1016/j.amc.2009.03.037

[13] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, 3rd ed. New York, NY: Wiley, 2006. https://doi.org/10.1002/0471787779

[14] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York, NY: Springer, 2006. https://doi.org/10.1007/978-0-387-40065-5

[15] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, 1963. https://doi.org/10.1080/01621459.1963.10500830

[16] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, 2nd ed. Philadelphia, PA: SIAM–MPS, 2014. https://doi.org/10.1137/1.9781611973433

[17] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton, NJ: Princeton Univ. Press, 2009. https://doi.org/10.1515/9781400831050

[18] D. Bertsimas, D. B. Brown, and C. Caramanis, "Theory and applications of robust optimization," *SIAM Review*, vol. 53, no. 3, pp. 464–501, 2011. https://doi.org/10.1137/080734510

[19] P. Mohajerin Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric," *Mathematical Programming*, vol. 171, pp. 115–166, 2018. https://doi.org/10.1007/s10107-017-1172-1

[20] D. Dubois and H. Prade, *Fuzzy Sets and Systems: Theory and Applications*. New York, NY: Academic Press, 1980. https://doi.org/10.2307/2273604
[21] G. J. Klir and B. Yuan, *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Upper Saddle River, NJ: Prentice Hall, 1995. https://doi.org/10.5555/202684

9 AUTHORS

Yogeesh N. is with the SR University, Warangal, Telangana, 506371, India; Research Fellow, INTI International University, 71800 Nilai, Negeri Sembilan, Malaysia (E-mail: yogeesh.n@ka.gov.in).

Markala Karthik is with the Department of Electrical and Electronics Engineering, SR University, Warangal, Telangana, 506371, India.

Asokan Vasudevan is with the Faculty of Business and Communications, INTI International University, Nilai 71800, Malaysia; International Institute of Management and Entrepreneurship, 220086, city Minsk, Slavinsky st.1/3, Republic of Belarus; Wekerle Business School, 1083 Budapest, Hungary (E-mail: asokan.vasudevan@newinti.edu.my).

Shankaralingappa B. M. is with the Department of Mathematics, Government First Grade College, India (E-mail: shankaralingappa.b@ka.gov.in).

Soon Eu Hui is with the Faculty of Business and Communications, INTI International University, Negeri Sembilan, Malaysia (E-mail: euhui.soon@newinti.edu.my).