International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Memetic Metaheuristic Initialisation of Neural Network Surrogates for Weight Optimisation of Steel Roof Trusses

Deepak Kulyal[1], V. K. Verma[2]

[1]Research Scholar, Department of Civil Engineering, College of Technology,
G.B. Pant University of Agriculture & Technology, Pantnagar, Uttarakhand, India
[2]Associate Professor, Department of Civil Engineering, College of Technology,
G.B. Pant University of Agriculture & Technology, Pantnagar, Uttarakhand, India
Corresponding author: deep.kulyal14@gmail.com

Abstract - Published studies disagree on whether initialising a neural network with a genetic algorithm (GA) or particle swarm optimisation (PSO) improves on ordinary backpropagation. This study tests the question under fair conditions for surrogates of steel roof truss behaviour. Hybrid multilayer perceptrons (MLP-GA, MLP-PSO) were compared with a plain MLP on four truss profiles (Fink, Howe, Pratt, Warren) designed to IS 800:2007. Each hybrid received the same backpropagation budget as its base model and a validation-based fitness. The key setting was a memetic fitness, in which each candidate received a short burst of backpropagation before being scored. With this setting, MLP-GA improved the mean testing R^2 by 0.005–0.012 and was more accurate on 27 of 32 profile–response combinations; MLP-PSO improved it by 0.004–0.008 and won 20. Neither hybrid lost on any profile. The surrogates were then used to select the lightest compliant design at 24 spans from held-out analysed designs. All 72 designs chosen by the three perceptron surrogates complied with the code when checked against their analysis results. The weight given up against the true lightest design averaged 2.4 % for MLP-GA, 2.9 % for MLP-PSO and 4.3 % for the plain MLP. Gains concentrated on the maximum utilisation ratio, the hardest response to predict. Hybridisation is therefore worthwhile mainly where the base network struggles.

Key Words: steel roof truss, surrogate-assisted optimisation, genetic algorithm, particle swarm optimisation, memetic algorithm, neural network.

1. INTRODUCTION

Metaheuristic optimisation of steel trusses requires thousands of structural analyses, one per candidate design. Replacing the analysis with a trained neural network surrogate removes most of this cost (Cho et al., 2007; Nourian et al., 2023). The surrogate's accuracy then decides the quality of the optimum.

A common way to improve a surrogate is to hybridise it with a metaheuristic. A GA or PSO searches the weight space for a good starting point, which backpropagation then refines (Ojha et al., 2017). Reported gains vary widely. Some studies report clear improvements from GA or PSO (Lee et al., 2025; Kao et al., 2025), while others find that combining global search with gradient training matters more than the metaheuristic itself (Kashyap et al., 2025).

Three weaknesses in study design can inflate or hide the gain. First, a hybrid that trains for longer than its baseline wins partly on extra training. Second, a global search scored on training error picks weights that already fit the training data, leaving little for backpropagation to improve. Third, scoring untrained weights asks where a candidate starts, not where training takes it.

This paper addresses all three for surrogates of steel roof truss behaviour. The contributions are:

1. A fair comparison of MLP, MLP-GA and MLP-PSO surrogates with an equal backpropagation budget and a validation-based fitness.
2. Evidence that a memetic fitness (Moscato, 1989), in which each candidate is briefly trained before scoring, is the setting that makes the hybrids reliable.
3. A surrogate-driven design search at 24 spans, verified against analysis results, which reports the weight given up to the surrogate and the code compliance of every selected design.

The work is part of a doctoral study (Kulyal, 2026). The database and the choice of the MLP over a multiplicative neuron model are reported in a companion paper (Kulyal and Verma, under review).

2. METHODOLOGY

2.1 Database and base surrogate

The surrogates were trained on a STAAD.Pro database of about 11,500 designs per profile for Fink, Howe (6–15 m spans), Pratt and Warren (15–40 m spans) trusses in Fe250 steel, with loads to IS 875 and checks to IS 800:2007. Passing and failing designs were both retained; its construction is described in the companion paper. Each profile was split 70:15:15, stratified on design status, and scaled to 0.1–0.9 using training statistics only.

The base surrogate is an MLP with nineteen inputs, one hidden layer of ten logistic units and a logistic output (211 parameters). Separate networks predict eight continuous responses, including total weight and maximum utilisation ratio, and one classifies design status. Training used mini-batch backpropagation with momentum 0.9 for 300 epochs, with early stopping on validation error.

2.2 Hybrid initialisation

In the hybrids, a GA or PSO searches the 211-dimensional weight vector for a starting point, and backpropagation then trains from it. Settings are given in Table 1.

Table -1: Global search settings

Setting	GA	PSO
Population / swarm	30	30
Generations / iterations	20	20
Operators	Tournament (size 3), BLX-α crossover (α = 0.5, rate 0.85), Gaussian mutation (rate 0.05, σ = 0.3), 2 elites	Constriction form, χ = 0.7298, $c_1 = c_2$ = 2.05, velocity limit ±0.5
Weight bounds	±3	±3

The PSO velocity update follows the constriction form of Clerc and Kennedy (2002):

$$v_i^{k+1} = \chi\left[v_i^k + c_1 r_1\left(p_i - x_i^k\right) + c_2 r_2\left(g - x_i^k\right)\right], x_i^{k+1} = x_i^k + v_i^{k+1}$$

2.3 Fair-comparison rules

Four rules make the comparison with the base MLP fair:

1. **Validation fitness.** Candidates are scored by mean squared error on the validation subset, not the training subset.
2. **Equal training budget.** Each hybrid receives the same 300 backpropagation epochs as the base MLP, so any gain comes from the starting point, not from longer training.
3. **Same initial distribution.** The population is drawn from the same random initialisation the base MLP uses.
4. **Memetic fitness.** Before scoring, each candidate is trained for four epochs on a random subsample of 1,500 training records. The search therefore compares candidates by where training takes them, not by where they start.

The hybrid does receive one advantage: the searched starting point costs 630 extra fitness evaluations per network. This cost is reported rather than hidden. Every configuration was trained from three fixed seeds and the best run on validation data was kept.

2.4 Surrogate-driven design search

To test the surrogates as design tools, each was used to pick the lightest compliant design at six spans per profile (24 spans in total). The candidate pool at each span was the set of held-out testing designs at that span. The surrogates never saw these designs in training, but every one had been analysed, so the true response of any design picked is known.

Each surrogate ranked the pool by predicted weight and returned the lightest design it predicted to pass on utilisation, deflection and design status. The true weight and utilisation of that design were then read from the analysis results. The true lightest compliant design in the same pool was used as the reference, so the weight given up to the surrogate can be stated as a percentage.

An acceptance margin below a utilisation of 1.0 was intended to absorb prediction error. It could not be applied here: in the industrial pools no design had a true utilisation below 0.97. The search therefore ran without its main safeguard, which makes it a demanding test.

3. RESULTS

3.1 Effect of the memetic fitness

In a first version, candidates were scored on the validation error of the untrained network. Gains over the base MLP were real but small. Adding the short inner training changed this: after the change, neither perceptron hybrid finished below the base model on any profile. Both searches improved quickly over the first generations and then flattened without collapsing onto a single solution, indicating that 20 generations were adequate.

3.2 Gain over the base model

Both hybrids improved the mean testing R^2 over the eight responses on all four profiles (Table 2). MLP-GA was the stronger, winning 27 of 32 profile–response combinations.

Table -2: Change in mean testing R^2 against the base MLP

Profile	MLP-GA − MLP	MLP-PSO − MLP
Fink	+0.0121	+0.0061
Howe	+0.0049	+0.0047
Pratt	+0.0094	+0.0045
Warren	+0.0114	+0.0075
Combinations won (of 32)	27	20

The gains were largest on the maximum utilisation ratio. For example, MLP-GA raised utilisation R^2 on the Warren truss from 0.8046 to 0.8493, while weight R^2 changed only from 0.9651 to 0.9665.

The same memetic scheme applied to a multiplicative neuron model gave inconsistent results, improving on some profiles and degrading on others. That model is not considered further; see the companion paper.

3.3 Design-status classification

The hybrids also lowered the false negative rate (FNR), the share of failing designs predicted to pass (Table 3). MLP-GA matched or improved on the base MLP on all four profiles.

Table -3: False negative rate for design status on the testing subset

Profile	MLP	MLP-GA	MLP-PSO
Fink	0.2125	0.1958	0.1917
Howe	0.0743	0.0743	0.0673
Pratt	0.0873	0.0744	0.0858
Warren	0.1050	0.0773	0.0802

3.4 Surrogate-driven design search

All three perceptron surrogates returned a design at every one of the 24 spans, and all 72 selected designs complied with IS 800:2007 when checked against their analysis results.

The weight given up against the true lightest compliant design averaged 2.4 % for MLP-GA, 2.9 % for MLP-PSO and 4.3 % for the base MLP. MLP-GA reduced this penalty on three of four profiles (Table 4). On the Warren truss it was slightly worse.

Table -4: Weight given up against the true lightest compliant design (%)

Profile	MLP	MLP-GA
Fink	8.7	4.6
Howe	2.7	0.3
Pratt	3.0	1.3
Warren	2.7	3.5

Table 5 shows the MLP-GA design at one mid-range span per profile. Across all 24 spans, the predicted weight of the selected design was within 9.3 % of its analysed weight.

Table -5: MLP-GA selected designs at mid-range spans

Profile	Span (m)	Weight pred. (kg)	Weight anal. (kg)	Error (%)	Util. pred.	Util. anal.
Fink	12	3,517	3,768	−6.6	0.978	0.983
Howe	12	3,275	3,493	−6.2	0.981	0.985
Pratt	30	34,746	32,431	+7.1	0.988	1.000
Warren	30	29,644	31,042	−4.5	0.953	0.997

On the Warren truss, the hybrid made the largest difference to the constraint. The mean absolute error in predicted utilisation of the selected designs fell from 0.357 for the base MLP to 0.056 for MLP-GA and 0.073 for MLP-PSO.

4. DISCUSSION

4.1 Where the gain comes from

Backpropagation and a population search do different work. Backpropagation performs tens of thousands of gradient steps and is very effective on a smooth error surface with one dominant basin. The GA or PSO performs only a few hundred evaluations, but it can compare distant regions of weight space and choose between them.

That ability matters only when the error surface is irregular. Total weight is a smooth sum of member contributions, so backpropagation finds a good fit from almost any start and the hybrids add little. Maximum utilisation steps whenever the governing member changes, giving an error surface with many local minima. This is where the hybrids gained most.

The memetic fitness is what lets the search exploit this. A starting point with low initial error can sit in a narrow valley that training cannot leave, while one with higher initial error can lead into a broad basin. Scoring after a short burst of training ranks candidates by where they lead.

This also offers an explanation for disagreement in the literature. The size of a reported hybrid gain depends heavily on how well the baseline was trained. A weakly trained baseline leaves a large gap for the metaheuristic to close, so the gain then reflects the baseline's weakness rather than the hybrid's strength.

4.2 Bias from optimising against a prediction

The predicted weight of the selected design was below its analysed weight at 17 of 24 spans, with a mean error of about −2.8 %. This is expected. Among designs of similar true weight, an optimiser minimising predicted weight prefers the one whose weight the network happens to under-predict. The effect is small here because weight is predicted accurately, but it will occur with any surrogate and should be allowed for.

The same effect acts on the constraint. On the Warren truss, MLP-GA under-predicted the utilisation of every selected design, by up to 0.104. All six designs still complied, but only because the analysed utilisations stayed at or below 1.0. A smooth curve of predicted optimum weight against span does not show this risk; only checking the selected designs against analysis does.

4.3 Limitations

The verification used designs already in the database. Designs proposed from the full section catalogue, which may lie outside that set, still need independent STAAD.Pro checks. The acceptance margin could not operate because the industrial pools lacked designs with spare capacity. Each result is the best of three runs, which is too few for a formal estimate of run-to-run spread. The study covers Fe250 steel and the stated span ranges only.

5. CONCLUSIONS

GA- and PSO-initialised MLP surrogates were compared with a plain MLP for steel roof trusses under an equal training budget and a validation-based memetic fitness. The main conclusions are:

1. Both hybrids improved on the base MLP on all four profiles. MLP-GA gained 0.005–0.012 in mean testing R^2 and won 27 of 32 profile–response combinations; MLP-PSO gained 0.004–0.008 and won 20.
2. The memetic fitness is the setting that makes hybridisation reliable. Scoring untrained weights gave small and inconsistent gains.
3. The gain concentrates on the hardest response, maximum utilisation, and is small for smooth responses such as weight. Hybridisation is worthwhile mainly where the base network struggles.
4. Used as a design tool at 24 spans, all 72 designs selected by the perceptron surrogates complied with IS 800:2007. MLP-GA gave up 2.4 % of weight on average against the true optimum, compared with 4.3 % for the base MLP.
5. Optimising against a prediction biases the selected design towards under-predicted weight and utilisation. Every design proposed by a surrogate should be confirmed by full analysis before use.

Future work should extend the database to designs with spare capacity so that an acceptance margin can operate, and verify designs proposed from the full section catalogue with independent analyses.

DECLARATIONS

Source of the work. This article is derived from the doctoral thesis of the first author (Kulyal, 2026). Results are reported from that study and are not published elsewhere in journal form.

Data availability. The parametric database is available from the corresponding author on reasonable request.

Conflict of interest. The authors declare no conflict of interest.

Funding. This research received no external funding.

Use of AI tools. An AI writing assistant was used to help draft and edit the manuscript text. All data, analyses and conclusions come from the authors' own work, and the authors take full responsibility for the content.

REFERENCES

1. Cho, Y., Xia, L., Hong, S., Kim, S. and Bae, J. 2007. Study of optimized steel truss design using neural network to resist lateral loads. *Key Engineering Materials*, 348–349: 405–408.
2. Clerc, M. and Kennedy, J. 2002. The particle swarm – explosion, stability, and convergence in a multidimensional complex space. *IEEE Transactions on Evolutionary Computation*, 6(1): 58–73.
3. IS 800:2007. *General Construction in Steel – Code of Practice*. Bureau of Indian Standards, New Delhi.
4. IS 875:1987. *Code of Practice for Design Loads (Other than Earthquake) for Buildings and Structures*, Parts 1–3. Bureau of Indian Standards, New Delhi.
5. Kao, C. Y., Hung, S. L. and Xu, P. J. 2025. Application of the artificial neural network and enhanced particle swarm optimization to model updating of structures. *Journal of Civil Structural Health Monitoring*, 15: 59–86.

6. Kashyap, G. S., Nafis, M. T. and Wazir, S. 2025. A study of hybrid and evolutionary metaheuristics for single hidden layer feedforward neural network architecture. *arXiv*.
7. Kennedy, J. and Eberhart, R. 1995. Particle swarm optimization. In: *Proceedings of ICNN'95 – International Conference on Neural Networks*, 4: 1942–1948.
8. Kulyal, D. 2026. *Development of a Hybrid Artificial Neural Network and Metaheuristic Optimization Framework for the Optimal Design of Steel Trusses*. Ph.D. thesis, G.B. Pant University of Agriculture & Technology, Pantnagar, India.
9. Kulyal, D. and Verma, V. K. Under review. Additive versus multiplicative neuron surrogates for predicting the structural response of steel roof trusses.
10. Lee, D., Noh, S., Kim, J. and Lee, S. 2025. Efficient hyperparameter optimization using metaheuristics for machine learning in truss steel structure cross-section prediction. *Buildings*, 15(8): 2791.
11. Moscato, P. 1989. *On Evolution, Search, Optimization, Genetic Algorithms and Martial Arts: Towards Memetic Algorithms*. Caltech Concurrent Computation Program, C3P Report 826.
12. Nourian, N., El-Badry, M. and Jamshidi, M. 2023. Design optimization of truss structures using a graph neural network-based surrogate model. *Algorithms*, 16(8): 380.
13. Ojha, V. K., Abraham, A. and Snášel, V. 2017. Metaheuristic design of feedforward neural networks: A review of two decades of research. *Engineering Applications of Artificial Intelligence*, 60: 97–116.