Academic Journal of Applied Sciences
ISSN: 3105-0549 (Print), ISSN: 3105-0557 (Online) | Vol. 4, No. 1, 2026

Multi-Objective Optimization of Manifold Microchannel Heat Sink Structures Using a Gaussian Process Surrogate and NSGA-II

Weifu Huang *, Wenyan Xue

The College of Mechanical and Energy Engineering, Guangdong Ocean University, Yangjiang, China

* **Corresponding author**: (Email: 1992480852@163.com)

Abstract: Thermal resistance, pressure drop and wall-temperature uniformity of a manifold microchannel heat sink are controlled by the same three structural parameters; they conflict, the mapping is strongly nonlinear, and a study of this kind rarely holds more than a few dozen samples, too few for a search that calls a solver at every step. The paper develops a general procedure from such a sample set to a design decision. Semi-empirical forms first fix the physical magnitudes of the three indices, and the calibration decides which of them an analytical expression can carry: pressure drop can (coefficient of determination 0.8994), thermal resistance cannot (0.6229), because two huge fitted coefficients cancel and leave the parameters unidentifiable. Surrogates are then compared on the same samples, Gaussian process regression against two tree ensembles, to establish the accuracy ceiling inside the sampled domain and to cut the cost of one objective evaluation from a simulation to a matrix product. NSGA-II searches the non-dominated front on that surrogate, TOPSIS returns an equal-weight compromise, and a Latin hypercube sweep over the weight simplex tests the stability of the recommendation against preference change. The procedure yields a design with no pin fins, a depth ratio of 4.48 and four rows, at 36.4% of the pressure drop of the thermal-resistance optimum for 4.4% more resistance, and that design stays optimal under 71 of 100 random weight vectors. The procedure does not depend on the specific geometry and transfers to other small-sample structural design problems of the same parameter–performance type.

Keywords: Gaussian process regression, NSGA-II, TOPSIS.

1. Introduction

Power density in integrated circuits keeps climbing under advanced packaging and wide-bandgap devices, and the heat flux per unit area now exceeds what air cooling can remove. Each increment in junction temperature degrades switching loss and thermal-fatigue life together, so heat removal has become a constraint on performance rather than an accessory of packaging. A straight-through microchannel forces all flow through one long passage, so heat transfer and flow resistance grow together and the outlet wall runs hotter than the inlet. Manifold structures split the passage into short segments through a distribution layer, so the coolant turns back before the boundary layer develops fully and both resistance and wall-temperature gradient fall. Pin fins inside the channel disrupt the boundary layer further and enlarge the heat-transfer area, at the price of form drag. Numerical work on pin-fin manifolds has reported the variation of Nusselt number and friction factor and found the effect of fin geometry on temperature uniformity to be non-monotonic [1], and embedded manifold cooling for GaN HEMT devices has demonstrated the thermal capability of such structures [2].

Design then becomes a trade-off among thermal resistance, which measures heat-removal capability, pressure drop, which fixes pumping power, and the ratio of wall-temperature standard deviation to mean, which measures thermal reliability. All three are controlled by the same three parameters, and moving one almost always buys one index at the cost of another. Two lines of work have formed. One fits machine-learning models to simulated data and lets the surrogate drive the search, chaining prediction and structural selection inside one framework [3] or ranking several regressors for microchannel heat sinks under small-sample conditions [4]. The other pairs response surfaces with evolutionary algorithms, applied to U-, I- and Z-type manifold layouts [5] and to porous minichannels [6]. Three gaps persist. Most studies reduce the objectives to two and demote temperature uniformity to a constraint, although it is what triggers hot spots and thermal-stress concentration. Weights are set equal or by experience, and the stability of the recommended design against a change of preference goes untested. Performance drift under manufacturing error is rarely quantified, although a surrogate trained on 84 points needs its response to perturbation examined.

Mechanism supplies the magnitudes, the surrogate the evaluation, the evolutionary algorithm the search, and uncertainty checks close the loop. Semi-empirical models first establish how each index varies and where the variation is non-monotonic; the three regressors are then compared on the same samples to fix the accuracy ceiling within the domain; NSGA-II [7] searches the three-objective front and TOPSIS [8] returns an equal-weight compromise, which Latin hypercube weight sweeping and Monte Carlo perturbation then test against preference and manufacture. Three results follow. Calibration identifies the failure of the thermal-resistance form in its structural term, so the switch to a surrogate rests on a criterion rather than an accuracy comparison. Placing optimization, weight robustness and perturbation robustness on one data chain makes "is this optimum worth recommending" answerable in numbers. Treating the row count as discrete, finally, bands the front and puts the compromise on the low-pressure-drop band.

2. Formalization of Parameters and Indices

2.1. Geometry and the Three Parameters

The system stacks a manifold distribution layer on a microchannel heat-transfer layer, with deionized water as coolant and aluminium nitride as substrate ($k_s = 200\mathrm{W/(m\cdot K)}$). The distribution layer splits the inlet flow among manifold units; inside each unit several rows of pin fins sit along the stream, and the fluid turns back before joining the outlet plenum. The variable geometry reduces to three parameters: the ratio of pin-fin diameter to channel width β, the ratio of manifold depth to microchannel depth γ, and the number of pin-fin rows N in one manifold unit. β changes both the heat-transfer area and the blockage of the cross-section, γ changes the thermal path length and hydraulic diameter, and N sets how often the boundary layer is rebuilt. Their geometric meanings do not overlap, which is what later gives the sensitivity ranking a physical reading. The domain is $\beta \in [0, 0.3]$, $\gamma \in [3, 4.5]$ and $N \in \{0, 2, 4, 6, 8, 10\}$, so the optimization is of continuous–discrete mixed type.

2.2. Indices, Samples and Assumptions

Thermal resistance is the junction-to-ambient temperature difference divided by dissipated power, pressure drop is the static pressure difference between inlet and outlet, and temperature non-uniformity is the wall-temperature standard deviation normalized by the mean:

$$T^* = \frac{\sigma_T}{\bar{T}} = \frac{1}{\bar{T}} \sqrt{\frac{1}{n} \sum_{i=1}^{n} \left(T_i - \bar{T}\right)^2} \tag{1}$$

All three are made dimensionless against a reference point and written R^*, ΔP^* and T^*, so they can enter one decision matrix. Each maps onto a separate physical channel — margin, pumping power, hot-spot location — and no pair substitutes for the third: high conductance with low pumping power may still leave a steep wall gradient.

The sample set holds 84 records pairing three parameters with three indices, with no missing values and values on the grid above (not a full factorial). Four premises underlie the modelling: steady flow with constant inlet mass flow, inlet temperature and outlet pressure, and an incompressible Newtonian coolant of constant properties; a uniform volumetric heat source reaching the wall through the substrate, radiation and natural convection neglected; sample values faithful to their combination under one dimensionless convention; and independent, normally distributed manufacturing deviations. The first two permit the series-resistance and additive-loss forms, the last two delimit where the surrogate and the perturbation analysis apply.

3. Three Semi-Empirical Mechanistic Models

No closed analytical solution survives the manifold–fin coupling, and dimensional analysis yields groups without coefficients, so each index takes a semi-empirical form: its structural term from heat-transfer and resistance relations, its coefficients from the samples.

3.1. Thermal Resistance Network

Heat passes from the source through the substrate to the channel wall and is then carried away by convection, giving two resistances in series:

$$R = R_{\mathrm{cond}} + R_{\mathrm{conv}} = \frac{\delta}{k_s A_s} + \frac{1}{h A_f} \tag{2}$$

With δ the substrate thickness, A_s the conduction area and A_f the convective area. The convective coefficient follows from $h = Nu\, k_f / D_h$ and a power-law correlation $Nu = C_0 Re^m Pr^n$, with C_0, m, n to be calibrated. Substituting $Re = \rho u D_h / \mu$ shows that β reaches R_{conv} through both A_f and D_h, that γ acts on the same pair through channel depth, and that N raises Nu by rebuilding the boundary layer. Writing the three paths in calibratable form:

$$R^*(\beta, \gamma, N) = R_0 + \frac{A}{1 + B\beta + C\beta^2} \cdot \frac{1}{\gamma^D} \cdot \exp(-EN) \tag{3}$$

Here $1/(1 + B\beta + C\beta^2)$ combines the fin effect on area with the blockage it introduces, γ^{-D} carries the depth ratio, and $\exp(-EN)$ expresses the diminishing return of extra rows; six coefficients are calibrated.

3.2. Pressure Drop

Total pressure drop superposes distributed friction and local form drag, $\Delta P = \Delta P_f + \Delta P_l$. The distributed term follows Darcy–Weisbach with a laminar friction factor $f = C_f / Re$, and the local term sums the contributions of fins, inlet, outlet and manifold turns, each weighted by its loss coefficient and dynamic head. Combining both and rendering the result dimensionless:

$$\Delta P^*(\beta, \gamma, N) = P_0 + A_p \frac{1 + B_p \beta^F}{C_p \gamma} \left(1 + D_p N^G\right) \tag{4}$$

The numerator is the form drag added by the fins, the γ in the denominator the dilution of resistance as the manifold deepens, and the exponents F, G let resistance grow faster than linearly; seven coefficients are calibrated.

3.3. Temperature Non-Uniformity

Non-uniformity arises from uneven flow distribution among manifold units and from spatial variation of the convective coefficient. Fins flatten temperature peaks but too many rows degrade the distribution in turn, and the depth ratio acts through the same channel, so the response carries quadratic and interaction terms:

$$T^*(\beta,\gamma,N) = T_0 - A_t\beta + B_t\beta^2 - \frac{C_t}{\gamma} + \frac{D_t}{\gamma^2} - E_tN + F_tN^2 + G_t\beta N \tag{5}$$

The quadratic terms $B_t\beta^2$ and F_tN^2 allow an interior minimum, $G_t\beta N$ describes the coupling of rib width with row count, and eight coefficients are calibrated, 21 in total.

3.4. Calibration and Parameter Trends

Calibration minimizes the residual sum of squares:

$$\min_{\boldsymbol{\theta}} \sum_{i=1}^{84} \left(y_i^{\exp} - f(\mathbf{x}_i;\boldsymbol{\theta}) \right)^2 \tag{6}$$

solved by Levenberg–Marquardt iteration, whose damping term switches between gradient descent and Gauss–Newton and so tolerates coefficients of very different scale. Results appear in Table 1.

Table 1. Calibration of the three semi-empirical models

Model	R^2	RMSE	Leading coefficients
Thermal resistance	0.6229	0.0068	$R_0 = -45.43$, $A = 46.14$, $B = 0.0025$
Pressure drop	0.8994	0.0088	$P_0 = 0.0367$, $A_p = 0.298$, $B_p = 117.2$
Temperature non-uniformity	0.7697	0.0093	$T_0 = 1.129$, $A_t = 0.600$, $B_t = 1.449$

Pressure drop is modelled best, with 90% of its variance explained. The low accuracy of the resistance form is not a calibration failure but a signal of mis-specification. The fitted $R_0 = -45.43$ and $A = 46.14$ are far larger than the index itself, which lies around 0.72–0.77, and they cancel almost completely, so the small difference comes from subtracting two large numbers and the parameters stay unidentifiable. The expression is also monotonic in β, whereas measured resistance falls and then rises, and a monotonic form cannot reproduce that.

Figures 1, 2 and 3 show how the three indices vary along fixed slices. With $\gamma = 3.75$ and $N = 5$ fixed, R^* falls from 0.760 to about 0.7454 and then climbs back, with a minimum near $\beta \approx 0.24$: too thin a rib leaves insufficient area, too thick a rib blocks the passage, and the mechanisms balance in between. On the same slice ΔP^* rises from 0.099 at $\beta = 0$ to 0.134 at $\beta = 0.30$, about 36%, with the slope steepening so that the penalty accelerates. In Figure 3 T^* is U-shaped in the depth ratio with its minimum inside the domain, so resistance and uniformity both hold interior minima while pressure drop carries no ceiling; selecting on resistance alone drives the design to the wide-rib end, where the pumping penalty never enters a single-objective criterion.

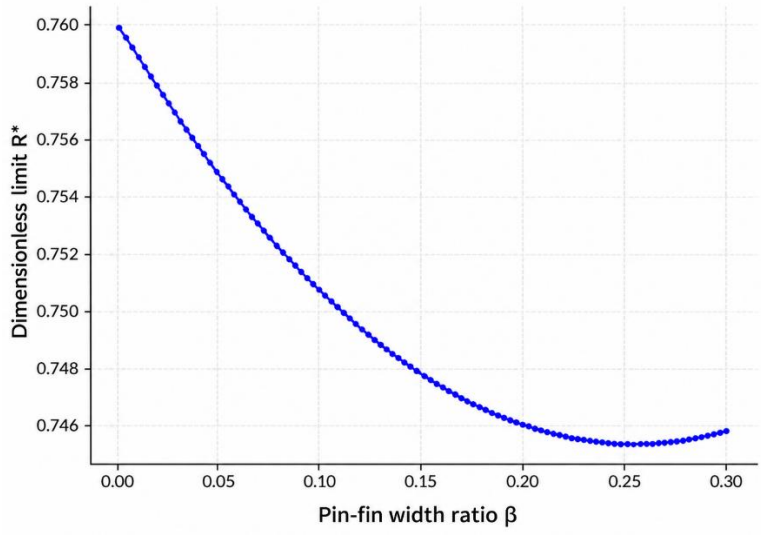


Figure 1. Effect of the pin-fin width ratio on dimensionless thermal resistance ($\gamma = 3.75$, $N = 5$)

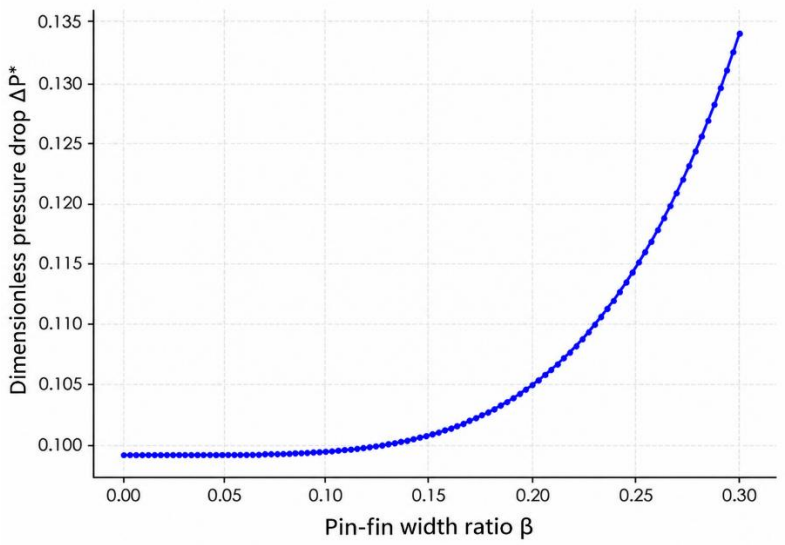


Figure 2. Effect of the pin-fin width ratio on dimensionless pressure drop ($\gamma = 3.75$, $N = 5$)

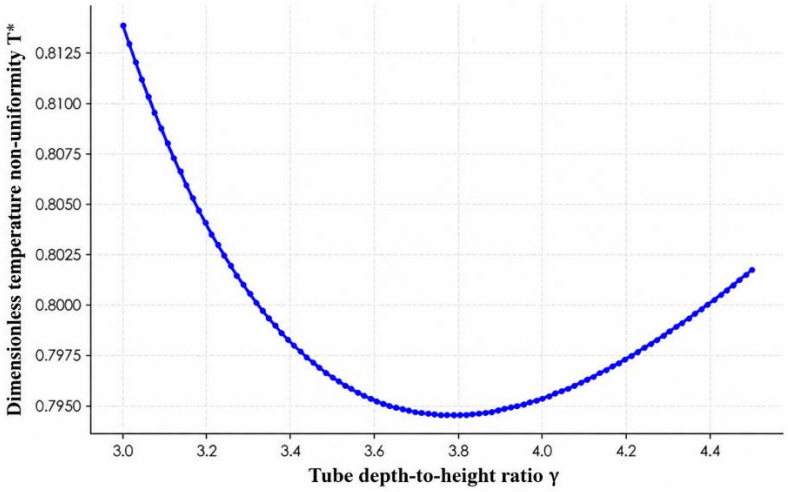


Figure 3. Effect of the manifold depth-to-height ratio on temperature non-uniformity

The three curves come from fixed slices: Figure 2's 0.099 holds at $\gamma = 3.75$, $N = 5$ and not elsewhere.

4. Surrogate Comparison

4.1. Protocol

Inputs are Z-score standardized to remove the scale gap between β and γ ; outputs are already dimensionless. Samples are split 80/20, with five-fold cross-validation inside the training set, and three measures are reported:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}, \qquad \mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}, \qquad \mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{7}$$

All three are needed because R^2 measures relative explanatory power only and inflates on outputs of small variance, while the outputs here span 0.72–0.77, 0.058–0.204 and 0.75–0.82.

4.2. Gaussian Process Regression

Gaussian process regression treats the unknown function as a sample path determined by a mean and a covariance function [9]. A constant mean and an isotropic radial basis kernel are used:

$$m(\mathbf{x}) = c, \qquad k(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right) \tag{8}$$

Where σ^2 is the signal variance and l the length scale. Given training data $\{X, \mathbf{y}\}$, the posterior at a new point stays Gaussian with

$$\mu_* = \mathbf{k}_*^{\mathrm{T}} \left(K + \alpha I\right)^{-1} \mathbf{y}, \qquad \sigma_*^2 = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^{\mathrm{T}} \left(K + \alpha I\right)^{-1} \mathbf{k}_* \tag{9}$$

K being the kernel matrix, $\mathbf{k}_*$ the test-to-train kernel vector and α the noise level; hyperparameters follow from maximizing the log marginal likelihood. The second term of the posterior variance is the reason for preferring GPR: predictive variance grows as a test point leaves the data, which supplies a readable confidence signal to the optimizer. With 84 samples no surrogate is reliable at the domain edge, so one reporting uncertainty serves better as an objective than one giving a point estimate.

4.3. Random Forest and XGBoost

Random forest averages $B = 100$ regression trees trained on bootstrap resamples, each split drawing $\sqrt{p}$ candidates from the three features, with depth capped at 8 and at least 2 samples per leaf. XGBoost [10] adds trees round by round against a loss regularized by leaf count and leaf-weight norm, at a learning rate of 0.05 and maximum depth 5. The first suppresses variance by bagging, the second lowers bias by boosting; how that shows on 84 samples is what the comparison answers.

4.4. Accuracy and Importance

Cross-validation results appear in Table 2 and Figure 4. GPR holds $R^2 \geq 0.9987$ on all three outputs with negligible between-fold spread, while the ensembles stay in the 0.83–0.95 band and XGBoost beats random forest on every output, so reducing bias is the more effective lever here. The gap in absolute terms is an order of magnitude: GPR errors lie at 10^{-5}–10^{-4}, XGBoost at 10^{-3}.

Table 2. Five-fold cross-validation of the three surrogates

Output	Model	Mean R^2	SD	Mean RMSE	Mean MAE
R^*	GPR	1.0000	0.0000	0.00006	0.00003
R^*	RF	0.8373	0.0600	0.00432	0.00338
R^*	XGBoost	0.8684	0.0518	0.00386	0.00280
ΔP^*	GPR	0.9987	0.0015	0.00065	0.00038
ΔP^*	RF	0.9063	0.0239	0.00819	0.00587
ΔP^*	XGBoost	0.9504	0.0260	0.00587	0.00425
T^*	GPR	0.9995	0.0002	0.00041	0.00023
T^*	RF	0.8362	0.0478	0.00749	0.00584
T^*	XGBoost	0.9330	0.0317	0.00475	0.00342

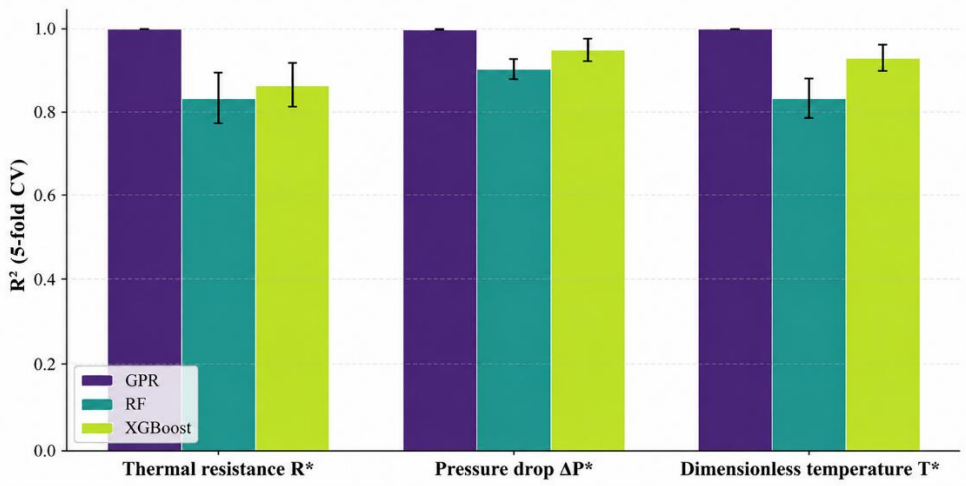


Figure 4. Five-fold cross-validated coefficients of determination for the three surrogates

On the independent test set GPR reaches R^2 of 1.0000, 0.9993 and 0.9996 with RMSE of 4.6×10^{-5}, 7.4×10^{-4} and 4.4×10^{-4}. Accuracy near unity comes from interpolation rather than learned physics: one dimensionless convention with smooth, noise-free outputs lets a radial basis kernel pass through the points almost exactly, and it is extrapolation that fails, the kernel falling back to the prior mean. The surrogate is therefore confined to the sampled domain, and so is the search.

Permutation importance from the random forest ranks the row count N first for resistance, and the depth ratio γ first for pressure drop and non-uniformity. Both readings are valid because they answer different questions: a slice slope is a local partial derivative with the other parameters pinned, whereas permutation importance is the error increase caused by shuffling one parameter across the space. Change the reading and the "most sensitive parameter" changes with it: the same data put γ first or β first.

5. NSGA-II Search and TOPSIS Decision

5.1. Mixed-Variable Model

The decision variables are the three parameters, the first two continuous and the third discrete. The three indices are minimized together under bound constraints alone:

$$\mathbf{F}(\mathbf{x}) = \left[R^*, \Delta P^*, T^*\right]^{\mathrm{T}} \quad \text{s.t. } 0 \leq \beta \leq 0.3,\ 3 \leq \gamma \leq 4.5,\ N \in \{0,2,4,6,8,10\} \tag{10}$$

No point minimizes all three: reducing β lowers pressure drop but raises resistance, adding rows does the reverse, and the effect of γ depends on the levels of the other two. Only a non-dominated set exists.

5.2. Non-Dominated Sorting and Diversity

NSGA-II [7] stratifies the population by dominance and keeps diversity by crowding distance, with boundary solutions assigned infinity so that the extremes survive:

$$d_i = \sum_{m=1}^{3} \frac{f_m^{(i+1)} - f_m^{(i-1)}}{f_m^{\max} - f_m^{\min}} \tag{11}$$

Tournament selection prefers the lower layer and, within a layer, the larger distance; offspring come from simulated binary crossover and polynomial mutation with distribution indices of 20, and N is rounded to the nearest admissible value, which sacrifices fine exploration but keeps offspring feasible. Settings are a population of 80, 150 generations, crossover probability 0.9 and mutation probability 0.2, fixed by convergence of the hypervolume indicator rather than by rule of thumb.

5.3. TOPSIS and Representative Designs

Closeness [8] ranks solutions by distance to the positive and negative ideals, after column-wise normalization and weighting by $v_{ij} = w_j r_{ij}$. All three objectives are cost-type:

$$A^+ = \{\min_i v_{i1}, \min_i v_{i2}, \min_i v_{i3}\}, \quad A^- = \{\max_i v_{i1}, \max_i v_{i2}, \max_i v_{i3}\} \tag{12}$$

With D_i^+ and D_i^- the distances to the two ideals, closeness is

$$C_i = \frac{D_i^-}{D_i^+ + D_i^-} \tag{13}$$

And equal weights $w_1 = w_2 = w_3 = 1/3$ are used. Vector normalization brings the objectives onto one scale but leaves their dispersion intact; pressure drop spans 0.058–0.204, the widest relative range, so its column dominates closeness under equal weights.

The hypervolume rises from 1.3×10^{-3} to 1.88×10^{-3} and flattens after about 85 generations, so the 150-generation budget is sufficient. Figure 5 shows the front of 80 non-dominated solutions, whose conspicuous feature is banding: with only six admissible values of N, the front is cut into sheets, each traced by continuous β and γ, with gaps between them. Tuning continuous parameters therefore buys little; what moves the front is the jump in row count. Four representative designs appear in Table 3.

Table 3. Performance of representative designs

Design	β	γ	N	R^*	ΔP^*	T^*
TOPSIS compromise	0.0000	4.4808	4	0.7521	0.0579	0.7801
R^* optimum	0.2316	3.0676	10	0.7203	0.1591	0.8196
ΔP^* optimum	0.0000	4.5000	4	0.7523	0.0575	0.7809
T^* optimum	0.0000	3.6082	4	0.7446	0.0805	0.7562

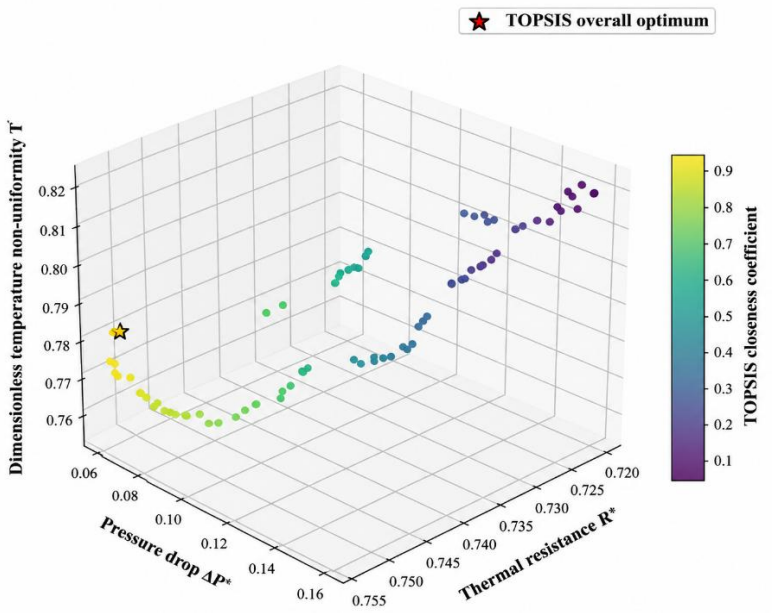


Figure 5. Three-dimensional distribution of the Pareto front (star: TOPSIS compromise)

The compromise is "no pin fins, deep manifold, moderate rows": $\beta = 0$ removes form drag, $\gamma = 4.48$ sits near the upper bound and dilutes resistance further, and $N = 4$ holds wall temperature even. Against the resistance optimum

it carries 36.4% of the pressure drop, 4.4% more resistance and 4.8% less non-uniformity, whereas the resistance optimum pushes resistance to 0.7203 and pays 0.1591, about 2.75 times the compromise. The pressure-drop optimum nearly coincides with the compromise (0.0575 against 0.0579, γ at the bound), which shows that equal weights let pressure drop dominate the decision. The temperature optimum lowers γ to 3.61 and improves uniformity to 0.7562 for 39.0% more pressure drop. Since the single-objective solutions sit at the extremities while the compromise sits at the pressure-drop end, equal weighting encodes an implicit "pumping power first" ordering that a recommendation should state.

6. Robust Design

6.1. Weight Sweeping

Equal weighting is only one preference. Weights are sampled over the simplex by Latin hypercube with 100 strata per dimension, rows normalized to sum to one, and the 2646 grid points ($21\times21\times6$) are evaluated once by the surrogate. One TOPSIS run per weight vector gives that vector's optimum, and each design's selection count is tallied. Frequency says "often best", not "reliably good", so high-frequency candidates are also scored by the standard deviation σ_i of their closeness, combined as

$$S_i = 0.5\cdot\frac{\mathrm{freq}_i}{\mathrm{freq}_{\max}} + 0.5\cdot\left(1-\frac{\sigma_i}{\sigma_{\max}}\right) \tag{14}$$

The halves ask whether a design wins under most preferences and whether it is ever much worse.

The 100 vectors produce 14 distinct optima. The five most frequent appear in Table 4.

Table 4. Most frequent optimal designs

Rank	β	γ	N	Frequency	R^*	ΔP^*	T^*
1	0.0000	4.5000	4	71/100	0.7519	0.0576	0.7806
2	0.0000	4.4250	4	9/100	0.7527	0.0590	0.7785
3	0.0000	4.2750	4	3/100	0.7531	0.0619	0.7735
4	0.0000	3.6000	4	3/100	0.7439	0.0807	0.7565
5	0.0000	4.3500	4	3/100	0.7531	0.0604	0.7761

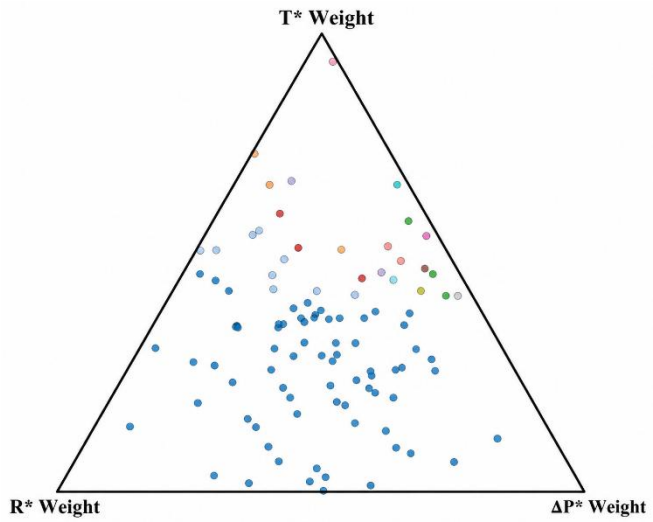


Figure 6. Ternary distribution of weights and the optimal design for each

All five share $\beta=0$ and $N=4$; only γ differs. The choice of rib width and row count is thus insensitive to preference, and what moves with preference is the manifold depth ratio — the quantity most easily controlled through etch depth once the other two are fixed by process. Figure 6 makes the same point geometrically: the points of the most frequent design occupy the lower-middle region and fan toward the ΔP^* vertex, while other colours cluster near the T^* vertex.

6.2. A Preference-Robust Design

Ranking by the combined score gives $\beta=0$, $\gamma=3.675$, $N=4$, compared with the compromise in Table 5.

Table 5. Comparison of power load forecasting of 403 line

Design	β	γ	N	R^*	ΔP^*	T^*
Equal-weight compromise	0.0000	4.4808	4	0.7521	0.0579	0.7801
Robust design	0.0000	3.6750	4	0.7453	0.0780	0.7567

Lowering γ from 4.48 to 3.675 cuts resistance by 0.9% and non-uniformity by 3.0%, and raises pressure drop by 34.7%. A compromise that equal weights rank below the pressure-drop optimum can be displaced once preference shifts, because it holds a clear margin along T^*; both designs have a place, the first where pumping power binds, the second where hot spots make the wall gradient critical. The value of the robust design lies not in being best under any single vector but in being least likely to collapse across the whole set.

6.3. Perturbation Stability

Normal perturbations at ±1%, ±3% and ±5% are applied to the compromise $(0, 4.48, 4)$ and the robust design $(0, 3.675, 4)$, with 100 Monte Carlo draws per level evaluated by the surrogate; since a relative perturbation of $\beta=0$ is meaningless, an absolute one with standard deviation 0.3 times the level is used. Statistics at the ±5%

level appear in Table 6, with the two distributions in Figure 7.

Table 6. Comparison of power load forecasting of 403 line

Design	Index	Mean	SD	CV	Range
Compromise	R^*	0.7534	0.0016	0.0021	[−0.0032, +0.0032]
Compromise	$\Delta ? P^*$	0.0617	0.0039	0.0630	[−0.0117, +0.0180]
Compromise	T^*	0.7819	0.0073	0.0093	[−0.0214, +0.0185]
Robust	R^*	0.7468	0.0038	0.0051	[−0.0069, +0.0074]
Robust	ΔP^*	0.0813	0.0079	0.0975	[−0.0166, +0.0202]
Robust	T^*	0.7649	0.0076	0.0099	[−0.0172, +0.0207]

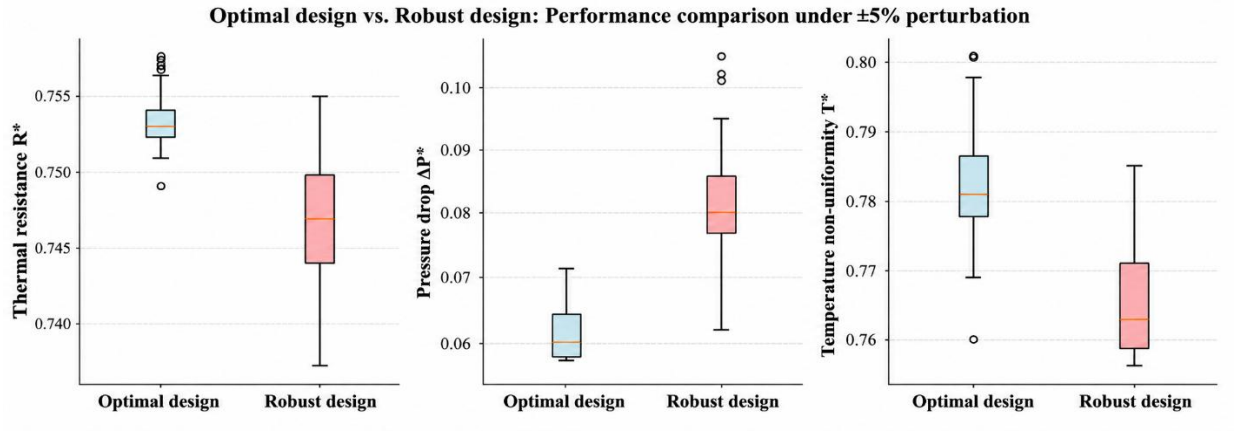


Figure 7. Performance distributions of the compromise and the robust design under ±5% perturbation

Pressure drop is the most sensitive, with a coefficient of variation of 6.3% for the compromise and 9.8% for the robust design; resistance is the least, at 0.21% and 0.51%; non-uniformity sits near 1%. The robust design is more sensitive because its γ is lower: with γ in the denominator, a smaller value amplifies the same relative perturbation. Lowering γ for uniformity also moves pressure drop onto a more sensitive operating point. Pressure-drop CV grows almost linearly with perturbation level, from 1.55% to 6.30% for the compromise and 1.91% to 9.75% for the robust design; the other two stay near a plateau.

The ranges are asymmetric: the compromise's mean pressure drop exceeds the nominal value by 6.6%, the upper excursion (+0.0180) exceeding the lower (−0.0117). This is not sampling error. Pressure drop is convex in γ, so a symmetric perturbation passed through a convex map lifts the mean and skews both bounds; a Monte Carlo mean is not the nominal design value.

The two dimensions give one rule: β and N are structural decisions that do not move, while γ carries preference and most of the variation. A design procedure should fix the first two, then choose γ within the application's balance between pressure drop and uniformity.

7. Conclusions

The three indices of a manifold microchannel heat sink respond to geometry in different shapes: resistance and non-uniformity non-monotonically, pressure drop monotonically and ever faster. Calibrating three semi-empirical forms on 84 dimensionless samples fixes the magnitudes, with pressure drop the most accurate ($R^2 = 0.8994$) and resistance the least (0.6229); in the latter two coefficients of huge magnitude cancel, so a single analytical form cannot carry the prediction and a data-driven surrogate becomes a necessity.

Gaussian process regression lifts all three outputs to $R^2 \geq 0.9987$ with absolute errors of 10^{-5}–10^{-4}, an order of magnitude better than the tree ensembles. The front stratifies along the discrete row count, so what relocates it is the jump in rows rather than fine tuning. Equal-weight TOPSIS returns $\beta = 0$, $\gamma = 4.48$, $N = 4$ at 36.4% of the pressure drop of the resistance optimum for 4.4% more resistance; because normalized pressure drop is the most dispersed, equal weighting implicitly orders "pumping power first", confirmed by the sweep that keeps the design optimal for 71 of 100 vectors. The robust design at $\gamma = 3.675$ trades 34.7% more pressure drop for 3.0% better uniformity and 0.9% lower resistance, and under ±5% perturbation its pressure-drop CV reaches 9.8% while resistance stays below 0.51%. Lock rib width and row count at the process stage, use depth ratio as the adjusting quantity, and monitor pressure drop. The method does not depend on the geometry at hand and transfers to small-sample design problems of the same parameter–performance type.

With 84 samples on a discrete grid the surrogate interpolates within the domain and nothing tests it outside; the optimization carries only variable bounds, with no feature-size, roughness, strength or cost constraint; uniform weights and normal perturbations were assumed, and real deviations may follow neither. Multi-output Gaussian processes or physics-constrained networks for extrapolation, constrained fronts including manufacturing cost, and a robust objective built from predictive variance are the natural next steps.

References

[1] Pan, Y. H., Zhao, R., Nian, Y. L., et al. (2023). Numerical study on heat transfer characteristics of a pin–fin staggered manifold microchannel heat sink. Applied Thermal Engineering, 219, 119436. https://doi.org/10.1016/j.applthermaleng.2022.119436

[2] Tang, W., Li, J., Lu, J., et al. (2023). Thermal management of GaN HEMT devices using subcooled flow boiling in an embedded manifold microchannel heat sink. Applied Thermal Engineering, 225, 120174. https://doi.org/10.1016/j.applthermaleng.2023.120174

[3] Yang, S., Zhao, X., Hu, W., et al. (2026). Machine-learning-based multi-objective optimization of a diamond manifold microchannel heat sink. Applied Thermal Engineering, 300, 131561. https://doi.org/10.1016/j.applthermaleng.2026.131561

[4] Sikirica, A., Grbčić, L., & Kranjčević, L. (2023). Machine learning based surrogate models for microchannel heat sink optimization. Applied Thermal Engineering, 222, 119917. https://doi.org/10.1016/j.applthermaleng.2022.119917

[5] Shanmugam, M., & Sirisha Maganti, L. (2023). Multi-objective optimization of parallel microchannel heat sink with inlet/outlet U, I, Z type manifold configuration by RSM and NSGA-II. International Journal of Heat and Mass Transfer, 201, 123641. https://doi.org/10.1016/j.ijheatmasstransfer.2022.123641

[6] Kumar, R., Zunaid, M., & Shyam Mishra, R. (2024). Multi-objective optimization of hydrothermal performance of a porous minichannel heat sink using RSM and NSGA-II algorithm. International Journal of Heat and Fluid Flow, 110, 109600. https://doi.org/10.1016/j.ijheatfluidflow.2024.109600

[7] Deb, K., Pratap, A., Agarwal, S., et al. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation, 6(2), 182–197. https://doi.org/10.1109/4235.996017

[8] Behzadian, M., Otaghsara, S. K., Yazdani, M., et al. (2012). A state-of-the art survey of TOPSIS applications. Expert Systems with Applications, 39(17), 13051–13069. https://doi.org/10.1016/j.eswa.2012.05.056

[9] Rasmussen, C. E., & Williams, C. K. I. (2006). Gaussian processes for machine learning. MIT Press. https://doi.org/10.7551/mitpress/3206.001.0001

[10] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785