Automated Constructal Design of Cold Plate Heat Exchangers with TPMS Unit Cells

Matei C. Ignuta-Ciuncanu & Ricardo F. Martinez-Botas

Imperial College London, UK, mi3618@ic.ac.uk

Abstract:

The growing need for efficient thermal management in electrified aviation systems, particularly for highperformance electronic components, presents a significant design challenge. This is further compounded by the common reductionist tendency to treat components in isolation. In this study, we extend the application of our generative bio-inspired design framework for cold plate heat exchangers, introducing hierarchical canopyto-canopy architectures built from triply periodic minimal surface (TPMS) unit cells. Unlike descriptive biomimicry or prescriptive parametric designs, the method is predictive as it autonomously identifies optimal configurations for arbitrary geometries by minimising entropy generation and adapting to changing performance demand. This design approach improves efficiency by providing increased design flexibility with a low number of degrees of freedom (DOFs, z=2-16), enhancing access to what flows (heat and fluid). The method integrates Constructal Theory with evolutionary algorithms and variational autoencoders (VAEs) to iteratively explore and optimise flow architectures. Through an evolutionary search process, the system discovers optimised designs that achieve an effective balance between competing thermal objectives, such as enhanced heat dissipation and minimised pressure drop. High-fidelity 3D computational fluid dynamics simulations validate the accuracy of the homogenized porous media model and elucidate the complex secondary flows (Dean vortices) responsible for enhanced mixing in Gyroid architectures. It is shown that evolutionary vascular networks reduce hotspot temperature non-uniformity by up to 50% while nearly halving the pumping power required in straight-channel parametric cold-plate designs. The approach benefits from rapid convergence using transfer learning to craft scalable architectures using a pre-trained constructal design model, reflecting nature's self-organising principles. This generative design methodology offers a new paradigm for developing highly efficient heat exchangers, with broad applicability in advanced thermal systems, including those for aviation and energy storage

Keywords:

Thermodynamics; Energy; ECOS Conference; Exergy; Sustainability.

1. Introduction

The electrification of aviation and the rapid increase in power density of modern electronic systems have placed unprecedented demands on thermal management. Heat fluxes in power modules now routinely reach several hundreds of W/cm2, rendering conventional air-cooling solutions insufficient and accelerating the transition toward liquid and two-phase cooling technologies [1-3]. Within this context, the geometric configuration of the coolant flow paths has emerged as a dominant factor governing both thermal performance and auxiliary power consumption.

Early advances in compact liquid cooling were enabled by microchannel heat sinks, which established a benchmark for high heat transfer coefficients in confined domains [1]. However, subsequent studies have shown that microchannels often suffer from excessive pressure drops and pronounced temperature nonuniformities, particularly when the imposed heat flux is spatially non-uniform [2].

In parallel, Constructal Theory has provided a unifying physical framework for understanding the emergence of efficient flow architectures in finite-size systems [4]. The theory states that systems persist by evolving configurations that facilitate easier access for the currents that flow through them, such as heat, mass, or fluid. Applied to thermal management, this principle predicts the natural emergence of hierarchical, tree-like, and vascular networks for area-to-point and point-to-area heat removal [5]. Importantly, Constructal Theory shifts the design question from prescribing a particular shape to identifying the configuration that minimises global imperfections, such as thermal resistance or flow dissipation, under given constraints. This perspective has motivated canopy-to-canopy (i.e., parallel area-to-area flow networks connecting distributed inlets and outlets through hierarchical headers) and branching flow architectures that outperform serpentine layouts in terms of hydraulic efficiency and temperature uniformity. Alongside constructal approaches, Topology Optimisation (TO) has become a powerful numerical tool for distributing material within a design domain to enhance heat transfer performance [6,7]. Density-based TO methods have produced highly complex, organic-like heat sink

geometries by solving large-scale optimisation problems directly coupled to governing transport equations. Despite their success, classical TO is computationally expensive, sensitive to mesh resolution, and prone to convergence toward local optima, particularly in multi-objective thermal–hydraulic problems [8] These challenges have motivated recent efforts to integrate machine learning with evolutionary optimisation to reduce computational cost while expanding the accessible design space [9].

A parallel development enabled by additive manufacturing is the practical realisation of Triply Periodic Minimal Surface (TPMS) structures, such as Gyroids, which offer high surface-area-to-volume ratios, smooth curvature, and favourable mechanical properties [10,11]. Experimental and numerical studies have demonstrated that TPMS-based heat exchangers can induce strong secondary flows that enhance mixing and convective heat transfer relative to conventional lattice or pin-fin geometries [12]. Their mathematical definition and geometric continuity make TPMS particularly compatible with generative design frameworks and hierarchical assemblies.

Despite these advances, a clear research gap remains. Existing parametric, constructal, and topology-optimised designs largely rely on prescribed geometric rules or high-dimensional optimisation spaces that limit adaptability to arbitrary heat-load distributions. Moreover, many optimisation strategies treat thermal and hydraulic objectives separately or rely on surrogate performance metrics that obscure the fundamental thermodynamic trade-offs at the system level.

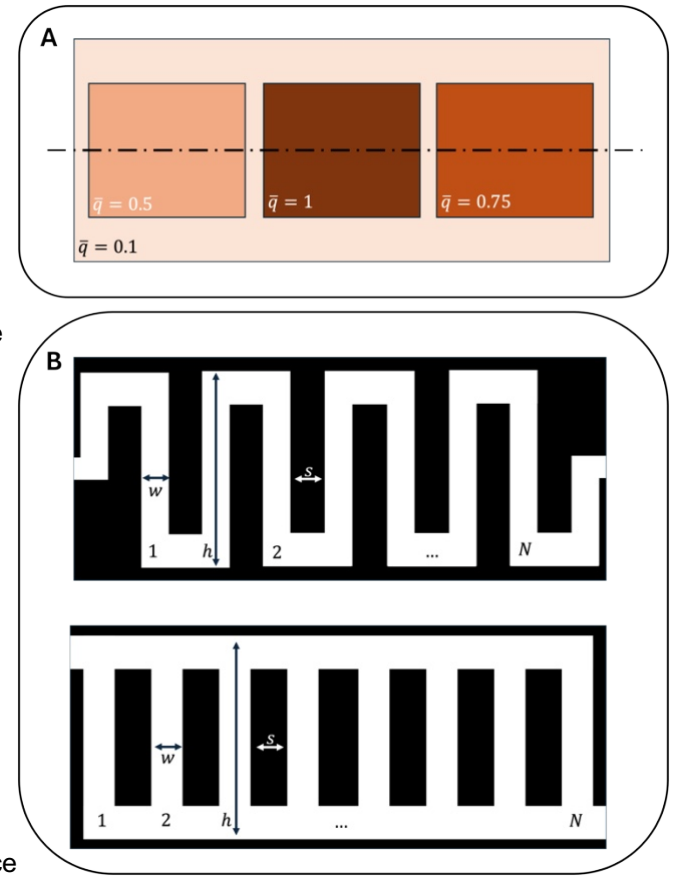

Figure 1. Representative parametric geometries (serpentine, canopy-to-canopy in B) illustrating the limitations of prescriptive cooling layouts under non-uniform thermal loads (A).

The present study addresses this gap by introducing a predictive, generative–evolutionary design framework for cold plate heat exchangers that unifies evolutionary design, machine learning, and thermodynamic optimisation. The approach encodes high-resolution two- and three-dimensional material distributions into a low-dimensional latent space using a variational autoencoder (VAE), with latent dimensionality $2 \leq z \leq 16$. This compact representation enables smooth interpolation between architectures and efficient exploration of a broad morphological design space. A multi-objective evolutionary algorithm (NSGA-II) operates directly within this latent space to evolve flow configurations that balance competing objectives, such as heat dissipation, hotspot mitigation, and pumping power. Each candidate design is evaluated using a conjugate heat-transfer solver, ensuring physics-based assessment without reliance on surrogate models.

2. Methodology

Parametric design remains the most widespread approach for cold-plate and heat-exchanger optimisation because it offers intuitive control over geometric descriptors such as channel width, aspect ratio, pitch, branching angle, or number of bifurcations [3] (Fig. 1B). Serpentine cooling is a prevalent baseline solution due to its manufacturability and reliable pressure-flow characteristics [16].

However, its single-path configuration enforces a fixed flow sequence, which inevitably introduces temperature gradients along the flow direction and leads to poor thermal uniformity under non-uniform heating. Hotspot mitigation is limited by the inability of the serpentine layout to reroute coolant preferentially toward thermally stressed components. Additionally, the long hydraulic path elevates pressure drop, reducing overall efficiency.

Recent developments within the constructal design literature demonstrate that canopy-to-canopy architectures overcome many of these limitations. Pioneering work by Bejan [17] provides analytical and numerical evidence that canopy arrangements reduce flow resistance and flatten temperature fields compared to serpentine channels. Gungor et al. [18] then showed that splitting and re-joining streams through multiple alternating canopies enhances access to the cooling potential of the fluid by redistributing flow more evenly across the hot surface. These findings support the adoption of hierarchical canopy-to-canopy networks as a superior baseline for thermal management systems where access to multiple local heat sources is essential.

In both cases, for a single channel, one may prescribe its aspect ratio, spacing to neighbours, curvature, or branching level; for a network of channels, the number of independent geometric choices grows combinatorially. Fractal and constructal branching parametrisations offer wider coverage of the morphological space by recursively defining successive orders of bifurcation but selecting the appropriate fractal order or

bifurcation depth *in silico* remains challenging without prior analytical study [19]. Each increase in fractal order multiplies the number of geometric degrees of freedom, making the search landscape irregular, multimodal, and difficult to navigate with classical sweeps.

These issues are aggravated when the cooling demand is highly localised. Parametric designs lack an intrinsic mechanism for morphology to *move* toward hotspots: instead, the designer must explicitly prescribe where and how channels branch or converge. This prescription-based nature restricts the ability of parametric methods to autonomously adapt (*self-organise*) to arbitrary heat maps [20].

2.1. Evolutionary Generative Design

To address the limitations of parametric prescriptions, we employ a generative–evolutionary design method grounded in latent-space modelling and evolutionary principles [8]. The generative agent maps high-dimensional shape information to a low-dimensional latent space using a convolutional variational autoencoder (VAE). This latent encoding provides a compact *genetic* representation of design candidates, enabling rapid, nonlocal shape variations while preserving physical realism and manufacturability [20].

In this context, the variational autoencoder (VAE) serves as a nonlinear dimensionality-reduction and generative modelling tool that learns a probabilistic mapping between complex geometric fields and a compact latent representation. The VAE consists of two coupled neural networks: an encoder, which compresses each high-dimensional material distribution into a low-dimensional latent vector by learning its underlying statistical features, and a decoder, which reconstructs physically plausible geometries from samples in the latent space.

From a thermal design perspective, the latent-space representation enables the evolutionary algorithm to reconfigure flow architectures in response to thermal demand without prescribing channel topology, branching rules, or unit-cell placement. Unlike parametric design, where geometric decisions are fixed *a priori*, or descriptive biomimicry, which imitates observed natural forms, the present approach is *predictive*: configurations emerge from physics-based evaluation of thermal and hydraulic performance.

Similar to topology optimization [14, 21], the method identifies architectures that improve heat spreading, hotspot mitigation, and pumping efficiency; however, the optimisation is not gradient-based and does not operate directly on a dense spatial field. Instead, evolutionary search in latent space enables nonlocal morphological changes, reduces sensitivity to local optima, and preserves manufacturable feature scales, while improving the hypervolume of the Pareto front by up to 13% [8]. The methodological pipeline is structured as follows:

- **Dataset and Transfer Learning**: The VAE is pretrained on a multimodal dataset comprising optimised pathways which encode general principles of heat diffusion, flow access, multi-scale branching, and self-similar transport patterns. By reusing the latent representations learned from bio-inspired heat-transport geometries, the model rapidly adapts to the thermal–hydraulic optimisation of cold-plate architectures with minimal retraining.
- **Latent representation:** Each candidate geometry is encoded by the VAE encoder as a latent vector $z \in \mathbb{R}^d$, where $2 \leq d \leq 16$ defines the effective number of latent degrees of freedom. Small perturbations in z therefore produce smooth, global variations in the decoded flow architecture, while suppressing high-frequency geometric noise. This continuous latent space provides a physically meaningful genetic representation that can be efficiently explored by the evolutionary algorithm.
- **Decoder Geometry Generation**: Candidate vascular networks are generated by decoding latent vectors, producing manufacturable, smooth channel morphologies without explicit parameterisation. Image-processing steps (thresholding, dilation, erosion) ensure continuity of flow channels.
- **Physics Evaluation**: Each decoded geometry is evaluated using a finite-element solver (FEniCS) for conjugate heat transfer performance [23]. This PDE-based evaluation enables direct optimisation of thermal–hydraulic performance without surrogate approximations.
- **Evolutionary Search**: The low-dimensional latent space is explored using NSGA-II, selected for its ability to maintain population diversity and resolve trade-offs between competing objectives. The multi-objective acumen is essential here because pumping power (pressure loss) and heat spreading must be optimised simultaneously, and the interactions between these objectives empirically lead to a non-convex search landscape [7], motivating a population-based optimiser.

This framework provides a predictive design tool where the configuration is not prescribed but emerges from the evolutionary search in latent space.

2.2. Hierarchical VAE and Multi-scale assembly

To capture the multi-scale nature of vascular networks, where two dendritic architectures can be matched canopy-to-canopy (Fig. 2B), we adopt the hierarchical VAE approach used for a metamaterial design task by our group at Imperial College [24].

First, a dedicated area-to-point network is required to distribute coolant from the inlet side toward the left-hand heat source, capturing the broad flow-access architecture. Second, a carrier unit cell provides the intermediary geometric layer that transports coolant across the domain. Third, a second area-to-point network is needed on the outlet side to collect and merge the distributed flow efficiently, completing the multi-scale conduit between source and sink (left to right in Fig. 2D).

This hierarchical assembly provides:

- **Coarse-to-fine control**, allowing the overall canopy-to-canopy encoding with reduced DOFs while permitting local tuning of channel curvature or thickness
- **Continuity constraints**, ensuring smooth transitions between neighbouring cells through filtered decoding and phase homogenisation
- **Design-space expansion**, enabling exploration of both global design (branch hierarchy, canopy arrangement) and local morphology (TPMS, as detailed in the next section).

This methodology echoes the multi-scale organisation of natural leaf venation, where global branching rules coexist with local anisotropies. By combining hierarchical VAE modelling with NSGA-II search, the resulting cold-plate designs achieve a balance between accessible morphological freedom and controlled manufacturability.

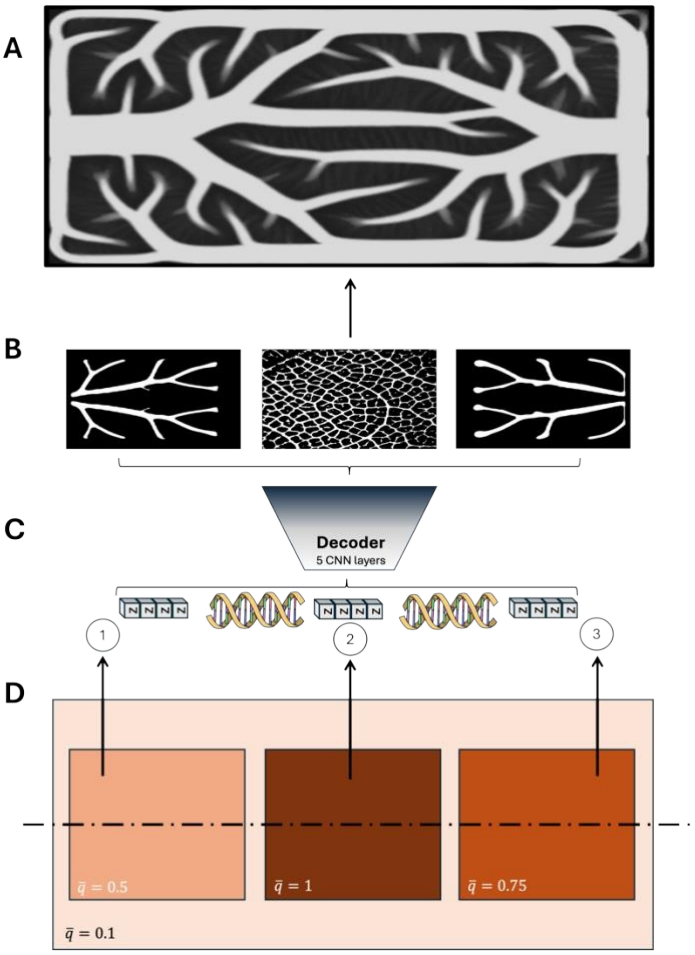

Figure 2. The Generative Hierarchical Assembly of the Canopy-to-Canopy Architecture. The design process maps the required local heat flux (D) to a low-dimensional latent vector (C). This vector is decoded to synthesize three regions: the inlet canopy (1), the carrier unit cell (2), and the outlet canopy (3) (B). The resulting full multi-scale vascular network (A) adapts its morphology to minimize entropy generation by balancing flow access and thermal contact.

2.3. Unit Cell Configuration

The effective heat transfer within the proposed hierarchical architecture relies on the integration of Triply Periodic Minimal Surfaces (TPMS) as the functional building blocks. Unlike traditional lattice trusses or pin-fin arrays, TPMS topologies are mathematically defined as surfaces with zero mean curvature at every point. This fundamental geometric property creates two interpenetrating fluid domains that are continuous in three dimensions, eliminating the stagnation zones often found in sharp-edged parametric geometries.

These structures are generated implicitly as the level set of a trivariate Fourier series approximation. The boundary between the solid wall and the fluid channel is defined where a scalar field function, $\psi(x)$, equals a constant isovalue *C*. Taking the Gyroid as the basic example, the governing implicit equation is:

$$\sin \pi x \cos \pi y + \sin \pi y \cos \pi z + \sin \pi z \cos \pi x = C \quad (1)$$

where x, y, z are spatial coordinates normalized by the unit cell periodicity.

Furthermore, the frequency of structural repetition is governed by the periodicity parameter ω ; this control is realized by modifying the trigonometric arguments to the general form $\sin \pi \omega x$ allowing the lattice density to be explicitly tuned. The complete library of unit cell topologies implemented in our generative map, including the Primitive, Diamond, and Lidinoid forms, is detailed in Fig. 3 with their characteristic equations.

The selection of TPMS unit cells is driven by their distinct impact on convective heat transfer mechanisms. The continuous, smooth curvature of the Gyroid induces complex secondary flows (Dean vortices) as the fluid navigates the tortuous channels. These secondary advective currents actively transport core fluid toward the

channel walls, disrupting the thermal boundary layer and enhancing the local Nusselt number without the severe pressure drop penalties associated with turbulent promoters.

Crucially, the implicit definition allows for seamless morphological adaptation. The isovalue *C* serves as a direct control parameter for the unit cell's porosity (volume fraction). By spatially modulating *C*, the generative algorithm can continuously tune the local hydraulic permeability and specific surface area. This capability enables the design to locally intensify heat transfer in hotspot regions (high surface area, high flow resistance) while opening up flow paths in cooler regions (low surface area, low flow resistance) to minimize global pumping power.

To ensure manufacturability of the TPMS-based unit cells, explicit bounds were imposed on the allowable porosity range. The local porosity was constrained to $\phi \in [0.50, 0.85]$, which prevents the formation of excessively thin ligaments at low porosity and avoids disconnected or mechanically fragile structures at high porosity. This range ensures minimum feature sizes compatible with metal additive manufacturing while preserving continuous fluid pathways and sufficient structural integrity [11].

Figure 3. Library of Triply Periodic Minimal Surface (TPMS) unit cells utilized in the generative design framework. The figure presents the iso-surface representation (left), the governing equations (center), and the resulting volumetric unit cell (right).

2.3. Flow and Thermal Solver: Porous Media Model

To enable the rapid, iterative evaluations required by the generative design framework, the thermo—hydraulic performance of candidate cold-plate geometries is assessed using a reduced-order Conjugate Heat Transfer (CHT) model based on the Darcy—Brinkman—Forchheimer (DBF) formulation [25]. The DBF approach provides an efficient representation of momentum and energy transport through complex porous structures without explicitly resolving the underlying micro-scale TPMS geometry.

Within the optimisation loop, the entire computational domain Ω—including inlet and outlet manifolds as well as the TPMS core—is treated as a single, continuous porous medium with spatially varying permeability and effective thermal properties. These properties are directly derived from the locally decoded TPMS morphology and porosity field. This formulation captures the essential three-dimensional flow resistance and heat-transfer behaviour of the gyroid-based structures while avoiding the prohibitive computational cost associated with meshing millions of individual TPMS features. The governing equations are solved using the Finite Element Method (FEM) within a custom implementation developed in the FEniCS framework [20]. The reduced-order porous-media solver is employed exclusively during the optimisation phase, allowing the evolutionary algorithm to explore a large design space efficiently and consistently under identical operating conditions.

Once the optimisation converges, selected Pareto-optimal designs are reconstructed as explicit three-dimensional geometries and subsequently analysed using high-fidelity, fully resolved 3D CFD simulations. These simulations are used solely for validation purposes and are not part of the optimisation process. This two-stage strategy ensures both computational tractability during design exploration and physical fidelity in the final performance assessment. The steady-state incompressible flow field, $\boldsymbol{u}$, and pressure field, p, are governed by the modified continuity and momentum equations, which collapse to the standard Navier-Stokes equations in the free-flow regions where the porosity $\phi \to 1$.

Continuity equation:

$$\nabla \cdot \boldsymbol{u} = 0 \tag{2}$$

Momentum (Darcy–Brinkman–Forchheimer) equation:

$$(\mathrm{u}\cdot\nabla)\mathrm{u} \;=\; -\nabla\,\mathrm{p} \;+\frac{\mu}{\phi}\,\nabla^{2}\mathrm{u} \;-\left(\frac{\mu}{\mathrm{K}(\phi)}\;\mathrm{u} \;+\frac{\mathrm{F}(\phi)}{\sqrt{\mathrm{K}(\phi)}}\;\rho\,|\mathrm{u}|\mathrm{u}\right) \tag{3}$$

where ρ is the fluid density, μ the dynamic viscosity, $\phi \in [0,1]$ the local porosity (design variable), and the three terms on the right-hand side denote the pressure gradient, Brinkman viscous diffusion and the Darcian/Forchheimer drag contributions respectively. The permeability and Forchheimer coefficient are interpolated as functions of permeability ϕ based on the correlations validated by Guillermo et al. [26].

Conjugate Energy Equation:

$$\rho\;\mathrm{C_p}\;(\mathrm{u}\cdot\nabla\,\mathrm{T}) = \nabla\cdot\;(\mathrm{k_{eff}}\,\nabla\,\mathrm{T}) \tag{4}$$

where C_p is the specific heat capacity and $\mathrm{k_{eff}}$ is the effective thermal conductivity, calculated based on a volume-average model: $k_{\mathrm{eff}} = k_f\,\phi + k_s\,(1-\phi)$ with k_f and k_s being the fluid and solid conductivities, respectively. The heat flux boundary condition, q''_{gen}, is applied to the base of the solid domain.

The critical link between the generative design process and the flow solver is the accurate prescription of the spatially varying resistance terms (K and F) for the DBF equation (Eq. 3). In this method, the local porosity $\phi(x)$ acts as the primary design variable and is the direct output of the VAE decoder. The VAE's latent vector z is mapped to a two-dimensional field of isovalues $C(x)$, which directly defines the local porosity $\phi(x)$ of the TPMS unit cell at every point x.

This methodology allows the evolutionary optimization algorithm to effectively tune the local flow resistance of the heat exchanger by varying the porosity field. For example, a lower porosity ϕ yields a lower local permeability K (increased pressure drop), allowing the generative framework to explore the optimal thermodynamic trade-off between heat transfer and pumping losses. The calibration of the flow solver with experimental flow data to account for Additive Manufacturing (AM) effects, a process common in similar topology optimization studies for pin-fin heat exchangers [27] is left as an opportunity for future work.

Thermal-hydraulic design objectives Effective thermal management requires optimizing both cooling efficiency and energy consumption. Three primary objectives are defined: one concerning the pumping power required to drive the system, and two concerning the temperature distribution.

Pumping Losses represent the energy dissipated due to the fluid flow through the system. These losses are characterized by an integral based on the velocity field and the Brinkman term, which accounts for both the viscous and inertial contributions to the flow resistance. The pumping losses are quantified as:

$$\Psi_f \;=\; \int_{\Omega}\left(\frac{1}{2}\mu\sum_{i,j}\left(\frac{\partial u_i}{\partial x_j}+\frac{\partial u_j}{\partial x_i}\right)+\sum_{i}\alpha(\phi)u_i\cdot u_i\right)d\Omega \tag{5}$$

where μ is the dynamic viscosity, $\alpha(x)$ is a term related to the flow resistance through the porous medium, and u is the velocity vector. Minimizing Ψ_f reduces auxiliary power consumption, at the cost of overall cooling performance.

Mean Temperature represents the overall thermal efficiency of the system, characterized by the average temperature across the surface. It is expressed mathematically as:

$$T_{\mathrm{mean}} \;=\frac{1}{\Omega}\int_{\Omega}\;T\;d\Omega \tag{6}$$

where Ω is the surface area and T is the local temperature. Lowering T_{mean} improves thermal management, leading to better performance and enhanced system longevity.

Hotspot Mitigation focuses on reducing localized high temperatures, which are critical for ensuring the longevity and efficiency of the system. The severity of hotspot formation is quantified using an Lp-norm of coefficient p=10, defined as:

$$T_{\mathrm{hot}} \;=\frac{1}{\Omega}\left(\int_{\Omega}\;(T-T_{\mathrm{mean}})^{p}\;d\Omega\right)^{1/p} \tag{6}$$

where T is the local temperature, T_{mean} is the mean temperature, and Ω is the surface area. Reducing T_{hot} minimizes localized temperature variations, ensuring better thermal uniformity.

3. Results and Discussion

3.1. Parametric Study

To establish a quantitative baseline for the generative design framework, a parametric study was conducted to compare the thermodynamic performance of traditional serpentine layouts against constructal canopy-to-canopy architectures [17]. This comparison validates the fundamental premise of Constructal Theory: that hierarchical "area-to-point" flow access yields superior efficiency compared to single-stream "point-to-point" configurations [28].

Computational Generation of Geometries

To ensure a rigorous comparison, the geometric configurations were generated procedurally using a custom Python script, imposing strict geometric constraints to isolate the effect of flow configuration. Both architectures were confined to a fixed domain footprint of height H=80 and width W=240 pixels. The primary design variable was the number of vertical heat-transfer branches, N, which was iterated through odd integer values in the range $N \in \{5, 7, \dots, 15\}$.

The generation rules are summarised as:

Serpentine Configuration: The vertical serpentine function generates a single continuous channel by connecting N vertical segments in series. The flow path alternates direction at each turn (connecting top-to-top or bottom-to-bottom), forcing the total mass flow to traverse the cumulative length of all branches. This results in a hydraulic path length that scales linearly with N.

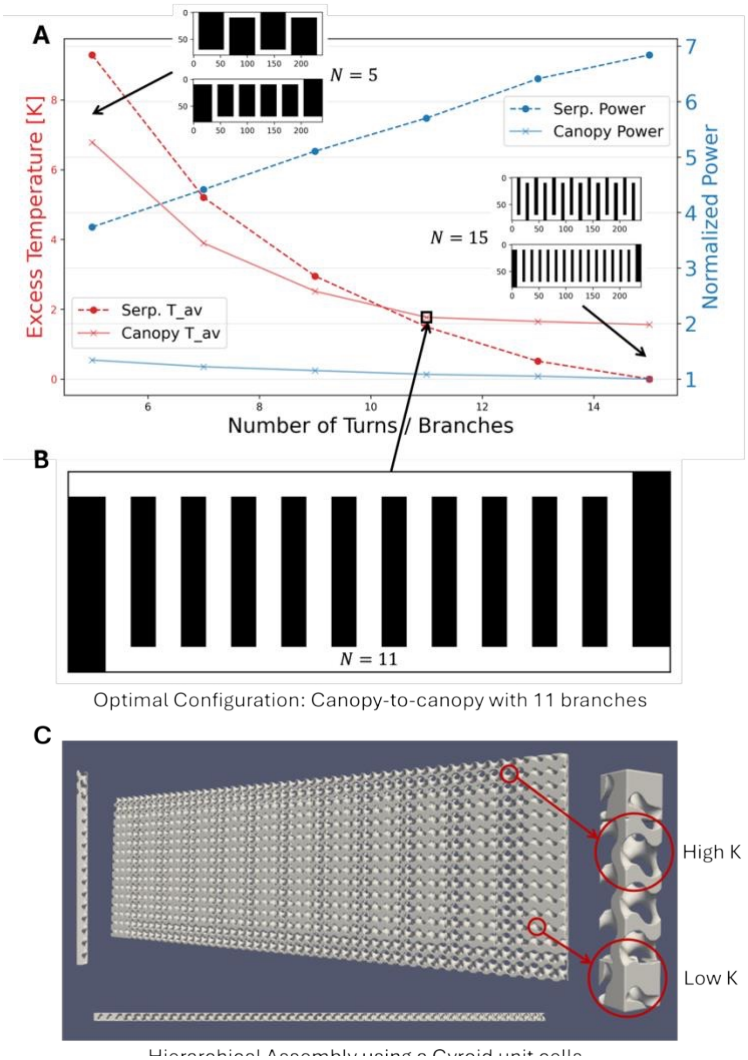

Figure 4. Parametric study results comparing serpentine and canopy-to-canopy architectures. (A) The trade-off between Excess Temperature (red) and Normalized Power (blue) as a function of branch count N. The canopy design maintains low power consumption while matching thermal performance within 2 K. (B) The optimal canopy configuration (N=11) identified by the study. (C) Visualization of the hierarchical assembly where the optimal canopy frame with Gyroid cells.

Canopy-to-Canopy Configuration: The canopy-to-canopy function connects the N vertical branches in parallel between two shared horizontal headers (inlet/outlet). This matches the elemental flow structures found in leaf venation [5], distributing the fluid simultaneously across the domain. The spacing between branches was calculated dynamically as $S = W / (N + 1)$ to ensure uniform coverage of the heat-generating surface.

Thermal and Hydraulic Performance Analysis

The conjugate heat transfer solver evaluated the steady-state performance of each candidate. The results, summarized in Fig. 4A, highlight the divergent scaling laws of the two architectures.

Thermal Performance (T_{mean}):

Both configurations exhibit a reduction in average excess temperature (T_{mean}) as the number of branches increases, driven by the increase in effective heat transfer surface area. However, the serpentine architecture (Fig. 4A, dashed red line) demonstrates superior cooling performance at lower branch counts compared to the canopy baseline (solid red line).

Hydraulic Power Consumption:

The most significant performance deviation is observed in the normalized pumping power required to drive the flow. The **Serpentine** Power (Fig. 4A, dashed blue line) exhibits a steep, linear increase as N rises. The series connection leads to a high aggregate flow resistance, as the pressure drop accumulates across every turn and channel segment. Conversely, the **Canopy Power** (solid blue line) remains nearly constant and up to 1 order of magnitude lower than the serpentine equivalent. By splitting the flow into N parallel streams, the canopy architecture reduces the local fluid velocity and the effective path length per fluid particle, drastically minimizing viscous dissipation.

Analysing the trade-off between thermal gain and hydraulic penalty identifies the canopy-to-canopy arrangement with $N = 11$ branches (Fig. 4B) as the optimal configuration. Beyond this, only marginal thermal reductions are outweighed by increased geometric complexity. This optimized macro-scale configuration is subsequently discretized using Gyroid unit cells to form a hierarchical porous architecture.

3.2. Macro-scale Configuration Evolution

The evolution of macro-scale configuration is driven by the need to optimize the complex interplay between thermal dissipation and fluid dynamics, framed by the overarching principle of entropy generation minimization. The optimisation loop shown in Figure 5 represents a fully automated evolutionary process in which candidate flow architectures are generated, evaluated, and evolved without manual intervention.

At each generation, a population of designs is assessed using the conjugate heat-transfer solver, and selection, crossover, and mutation are applied in latent space via the **NSGA-II** algorithm. Convergence of the evolutionary process is defined quantitatively by the stabilization of the Pareto front: the optimisation is terminated when the hypervolume of the Pareto set changes by less than 1% over successive generations, indicating that further improvements in the trade-off between objectives have saturated. The evolutionary loop is capped at a maximum of $N = 100$ generations, after which the optimisation is stopped even if the hypervolume criterion is not met.

The loop therefore does not involve post-processing or manual adjustment of intermediate designs. Instead, the flow physics is re-solved only for newly generated candidate architectures produced by the evolutionary operators, ensuring a consistent and objective exploration of the design space while progressively improving thermal performance and reducing pressure losses.

All results presented in this section are reported relative to a fixed serpentine cold-plate configuration, which

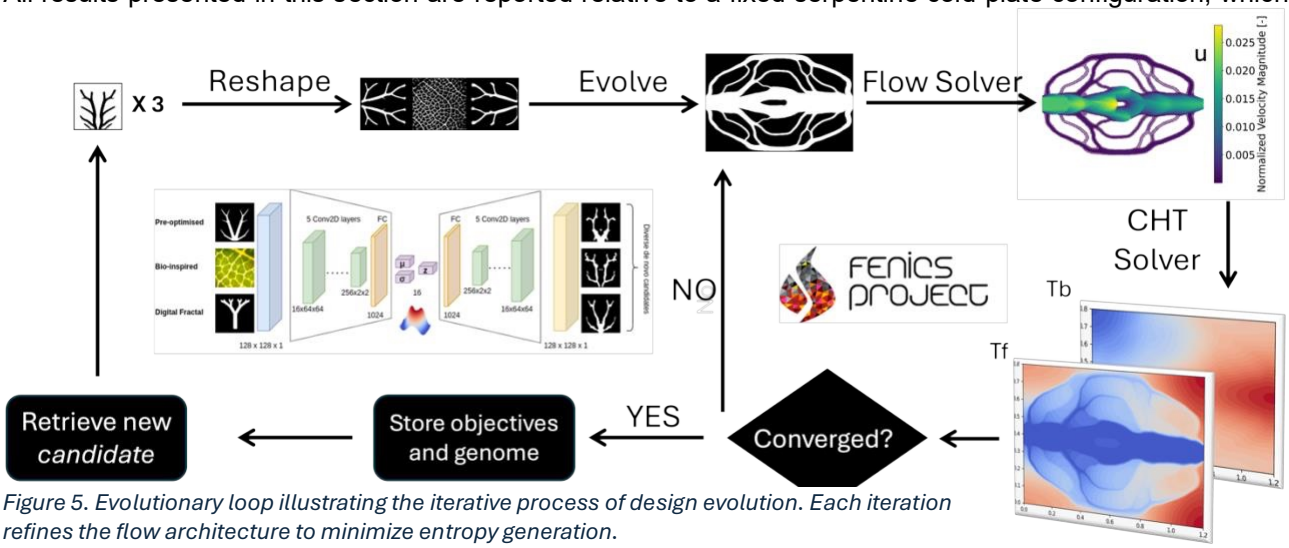

Figure 5. Evolutionary loop illustrating the iterative process of design evolution. Each iteration refines the flow architecture to minimize entropy generation.

is adopted as the reference (baseline) design. The serpentine layout represents a widely used industrial standard for liquid-cooled electronics and provides a consistent benchmark against which improvements in thermal performance and pumping power can be quantified. Performance gains reported throughout this section—including mean temperature reduction, hotspot mitigation, and pumping power savings—are therefore normalized with respect to this serpentine baseline, as summarized later in Table 1.

The evolutionary loop begins with configurations assembled from the generative-constructal design library published in [20]. These candidate designs are evaluated based on their thermal and hydraulic performance, with the design space explored iteratively using the **NSGA-II** algorithm. Over successive generations, the design approaches converge toward a Pareto optimal set of solutions, with the trade-offs between competing objectives becoming clearer.

As illustrated in Figure 6, the Pareto front represents the optimal trade-off between conflicting objectives. Here, we observe that designs such as the "Min Hotspot" (Design D) offer significant improvements in hotspot temperature reduction, while configurations like "Min Average T" (Design E) reduce the overall temperature gradient. This allows for a more nuanced understanding of how configurations perform under varying thermal conditions and highlights the flexibility of the evolutionary process to adapt to different design demands.

The evolutionary process emphasizes the impact of flow configuration (*design*, *drawing*) in thermal management. For instance, in Figure 7, we see how the generative design approach allows for spatially adaptive thermal management, where the fluid flow paths are reconfigured to more effectively target hotspots. This dynamic adjustment in geometry is in stark contrast to traditional fixed-flow designs like the serpentine configuration, which cannot adjust based on local heat generation.

The results from the evolutionary design process, as shown in Table 1, highlight the significant improvements achievable over conventional designs. In particular, Design D (Min Hotspot) achieves a remarkable +84% improvement in hotspot temperature, emphasizing the efficacy of evolutionary vascular networks in reducing thermal non-uniformities. Similarly, Design E achieves a 65% reduction in power while reducing the excess temperature across the surface by 80%, indicating not only a more efficient thermal management solution but also a reduction in energy costs for cooling. In contrast, configurations like Design C (Min Power) demonstrate a more drastic reduction in pumping power (up to 90%), although this comes at the expense of an increase in hotspot temperature (+35%). This result reflects the complex trade-offs inherent in the design space, where prioritizing one objective often leads to degradation in another, accentuating the necessity of a balanced approach in system optimization.

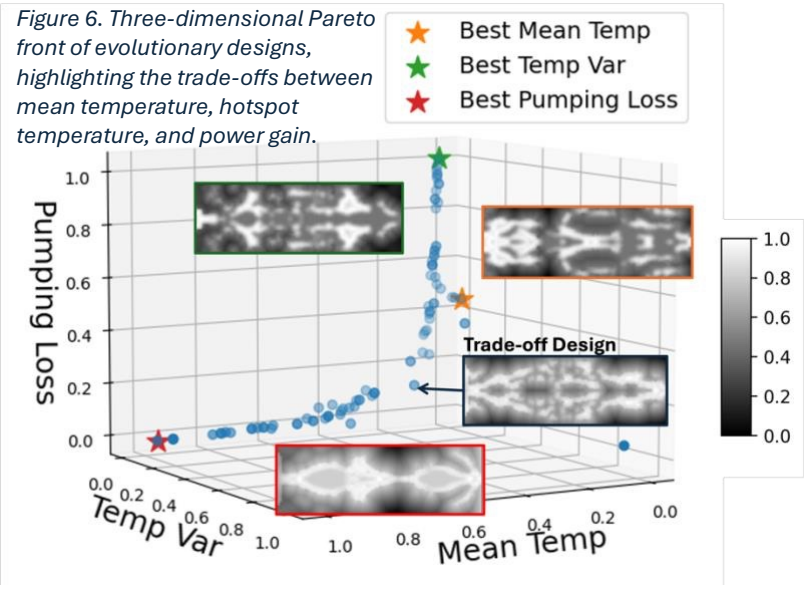

Figure 6. Three-dimensional Pareto front of evolutionary designs, highlighting the trade-offs between mean temperature, hotspot temperature, and power gain.

Meanwhile, Design F (Trade-off) demonstrates an optimal balance, with 80% improvement in mean temperature, 45% reduction in hotspot temperature, and 80% improvement in power. This trade-off design illustrates the effectiveness of evolutionary methods in navigating the multi-objective landscape, achieving significant improvements across multiple thermal and fluidic metrics.

The evolving designs offer valuable insights into the future of cold plate heat exchangers, particularly in the context of advanced aerospace applications and energy systems where heat fluxes are high, and localized cooling is crucial. The capacity for the generative algorithm to discover *self-organizing*, multi-scale architectures could play a pivotal role in optimizing thermal performance across a wide range of applications.

3.2. Multi-scale 3D results

To validate the accuracy of the porous media approximation used in the generative loop, high-fidelity three-dimensional Computational Fluid Dynamics (CFD) simulations were conducted. While the homogenized model efficiently guides the evolutionary search, full-scale CFD is necessary to capture the complex secondary flows and local pressure gradients induced by the intricate TPMS architectures.

Design	Mean Temp. (R_T)	Hotspot Temp. (R_H)	Pumping Power (R_P)
Serpentine (A)	1.000	1.000	1.000
Canopy (B)	0.617	0.783	0.356
Min Power (C)	1.784	0.687	0.070
Min Hotspot (D)	0.256	0.160	1.345
Min Average T (E)	0.197	0.617	0.375
Trade-off (F)	0.209	0.532	0.195

TABLE 1. Performance ratios of different designs compared to the Serpentine baseline. The ratio $\mathbf{R} < 1$ indicates a reduction (improvement) in the metric, while $\mathbf{R} > 1$ indicates an increase (degradation). $\mathbf{R} = 1 - (\text{Gain}/100)$.

A total of five distinct configurations were evaluated to rigorously benchmark hydraulic performance. A baseline case featuring a standard canopy architecture populated with Primitive unit cells was established to serve as a reference. Subsequently, four vascular flow networks—architecturally identical at the macro-scale but differentiated by their micro-scale unit cell topology—were simulated. The TPMS unit cells investigated include Primitive, Diamond, Lidinoid, and Gyroid structures, representing a broad spectrum of surface-area-to-volume ratios and tortuosity characteristics.

The simulations were performed using ANSYS Fluent [29], employing a pressure-based coupled solver. The computational domain discretized the full vascular geometry, requiring high-resolution unstructured polyhedral meshes to accurately resolve the boundary layers within the micro-channels. A uniform velocity inlet boundary condition was imposed, corresponding to a Reynolds number of $Re = 200$, a regime representative of electronics liquid cooling [18]. The working fluid was modelled as water with constant density and viscosity. Turbulence was modelled using the $k-\omega$ SST model to account for the transitional flow regimes expected within the tortuous TPMS lattices. Convergence was rigorously monitored, with residuals for continuity, momentum and turbulence equations required to drop below 10^{-5} to ensure numerical stability and accuracy.

To confirm the accuracy of the homogenized porous media model employed in the generative loop, its predictions were benchmarked against high-fidelity three-dimensional Computational Fluid Dynamics (CFD) simulations. As illustrated in Figure 8, the agreement between the two methods is excellent. The macroscopic velocity streamlines and pressure distributions predicted by the FEA porous model align closely with the volume-averaged fields obtained from the full 3D CFD.

Crucially, the pressure drop discrepancies remained within acceptable engineering limits (<15%) for all gyroid-like structures tested (excluding the Primitive), verifying that the Darcy-Brinkman-Forchheimer formulation adequately captures the bulk hydraulic resistance of the structure.

The observed low discrepancy for the Gyroid unit cell confirms its fitness as the base geometry used for the porous media coefficient calibration, suggesting that enhanced accuracy for multi-TPMS designs could be achieved by fitting individual calibration coefficients for each TPMS.

Moreover, the porous media estimate is always conservative, as this approach underestimates the maximum velocity achievable in less densely-packed regions. Notably, similar trends were observed by Banthiya et al. [27] for pin-fin heat exchangers with variable permeability.

The most significant advantage of the porous FEA approach lies in its computational efficiency. While a single high-fidelity CFD simulation of the complex vascular geometry required over 4000 seconds to converge, the homogenized FEA model completed the same analysis in less than 10 seconds—a speed-up of over two orders of magnitude. This drastic reduction in computational overhead is transformative, enabling the evolutionary algorithm to evaluate thousands of candidate designs in the time it would take to validate just one using traditional CFD. This efficiency justifies the use of the porous model as the primary engine for design exploration, reserving full 3D CFD for the validation of selected optimal candidates.

The primary metric for evaluation was the total pressure head required to drive the flow, a direct proxy for pumping power. The results in Table 2 reveal distinct hydraulic behaviours governed by the micro-scale features of each TPMS. The streamline analysis (Figure 9) demonstrates that while all configurations successfully distribute flow throughout the vascular network, the local vorticity and mixing intensity vary significantly. Structures like the Gyroid and Lidinoid induce strong, coherent secondary flows (Dean vortices) that enhance fluid mixing perpendicular to the main flow direction, potentially augmenting heat transfer at the cost of increased pressure drop. Conversely, the Primitive and Diamond structures exhibit more stratified flow patterns with lower hydraulic resistance but reduced mixing efficacy.

Beyond the hydraulic behaviour, the results also highlight the strong correlation between unit-cell topology and effective surface area, which is a key predictor of convective heat-transfer enhancement. As shown in Table 2 increased geometric complexity (L, G) exhibit substantially higher specific surface areas, up to 68% above the Primitive reference. This increased interfacial area provides a larger solid–fluid contact region over which heat can be exchanged, and, when combined with the intensified secondary flows observed in Figure 9, suggests a meaningful enhancement in thermal performance. However, this benefit comes with the trade-off of higher flow resistance, indicating that optimal designs must balance the competing effects of increased heat-transfer area and pressure drop.

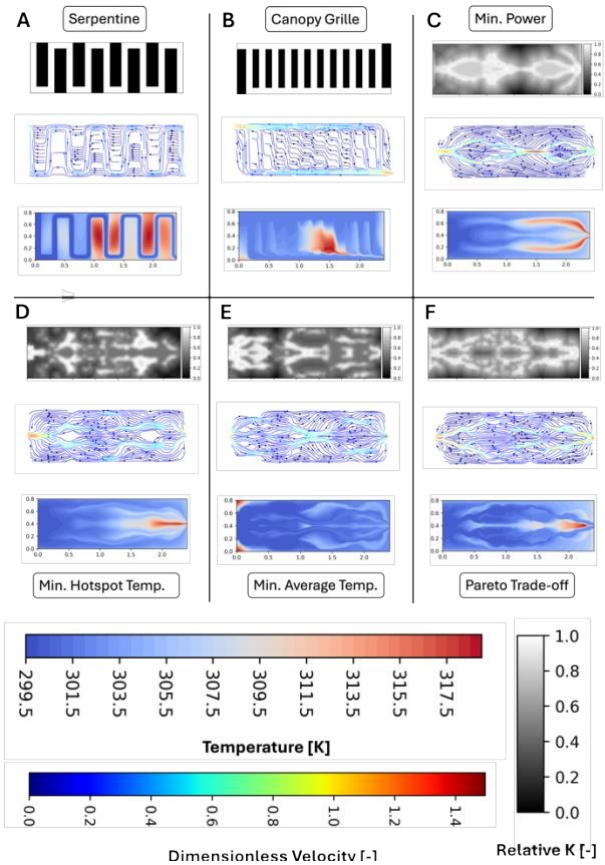

Figure 7. Heat map showing the multi-scale adaptive flow architecture of the generative design process. The fluid flow paths are reconfigured to minimize thermal gradients, effectively targeting hotspots while optimizing global heat dissipation.

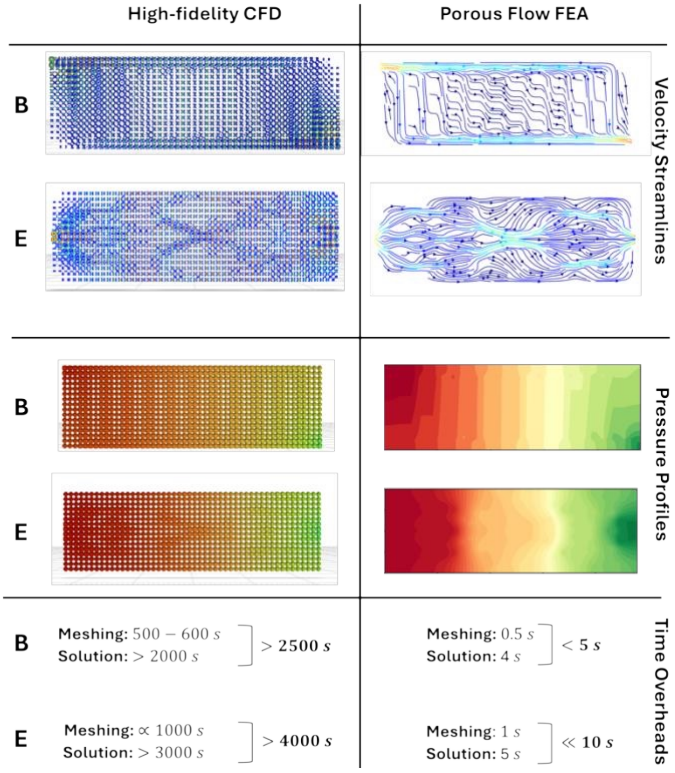

Figure 8. Comparative validation of flow physics between High-fidelity 3D CFD (left) and the Porous Flow model (right) for designs B and E. Top row: Velocity streamlines indicate that the homogenized porous model accurately captures the macroscopic flow distribution. Middle row: Pressure contours show excellent agreement in the hydraulic gradient across the domain. Bottom row: Time overheads.

4. Conclusion

This study successfully extends a generative bio-inspired design framework for cold plate heat exchangers, integrating hierarchical canopy-to-canopy architectures with Triply Periodic Minimal Surface (TPMS) unit cells. By combining Constructal Theory with evolutionary algorithms (NSGA-II) and Variational Autoencoders, the methodology moves beyond prescriptive parametric designs to autonomously identify configurations that minimize entropy generation.

The key findings from the hierarchical optimization process are summarized as follows:

- **Parametric Baseline Comparison**: The constructal canopy-to-canopy architecture demonstrated a hydraulic power consumption up to one order of magnitude lower than traditional serpentine layouts. The study identified an optimal macro-scale straight-channel configuration of $N = 11$ parallel branches, which balances thermal gain against geometric complexity.
- **Evolutionary Macro-Scale Performance**: The generative algorithm successfully navigated the multi-objective landscape to outperform the serpentine baseline significantly. The Trade-off design yielded an 80% improvement in mean temperature and pumping power while reducing hotspot temperature by 45%.
- **Multi-Scale Validation**: The homogenized porous media model used for optimization was validated against high-fidelity 3D CFD, showing pressure drop discrepancies of less than 5% compared to the gyroid TPMS. This approach reduced computational time from over 4000 seconds per simulation to under 10 seconds. Furthermore, the Gyroid unit cells provided up to 68% higher specific surface area compared to Primitive cells, while also inducing secondary Dean vortices that enhance mixing.

Configuration	Pressure Drop [kPa]	Relative Surface Area
Porous FEA (Baseline)	7.167	N/A
Primitive (P)	5.842	1.00 (Ref)
Diamond (D)	6.150	1.32
Lidinoid (IWP)	6.685	1.60
Gyroid (G)	6.920	1.68

TABLE 2. Comparison of total hydraulic pressure drop across the vascular network for different TPMS unit cells at $Re = 200$.

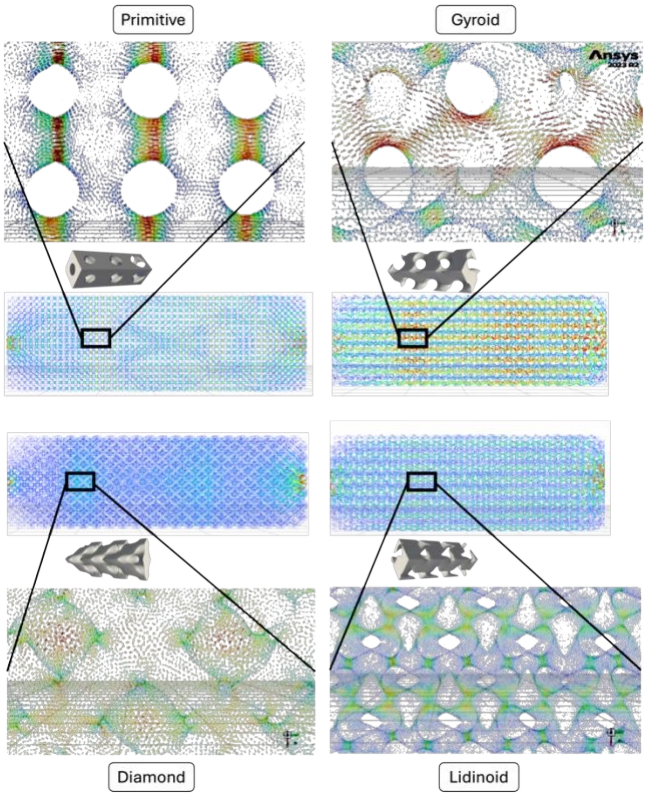

Figure 9. Comparison of velocity streamlines within the vascular network for four TPMS unit-cell topologies (Primitive, Gyroid, Diamond, and Lidinoid).

Ultimately, this work illustrates that thermal optimization is not achieved by imposing a fixed shape, but by liberating the architecture to change configuration to improve access to the currents that flow through it [30]. The design evolves by self-organizing the flow channels from the vascular macro-scale to the gyroid micro-scale to reduce thermodynamic imperfections.

References

[1] Tuckerman, D. B., and Pease, R. F. W., 1981. “High-performance heat sinking for VLSI”. IEEE Electron Device Letters, 2(5), pp. 126–129.

[2] Mudawar, I., 2013. “Recent advances in high-flux, two-phase thermal management”. Journal of Thermal Science and Engineering Applications, 5(2).

[3] Bejan, A., and Errera, M. R., 2015. “Technology evolution, from the constructal law: heat transfer designs”. International Journal of Energy Research, 39, 6, pp. 919–928.

[4] Bejan, A., 1997. Advanced Engineering Thermodynamics, 2nd ed. John Wiley & Sons.

[5] Bejan, A., 2000. Shape and Structure, from Engineering to Nature. Cambridge University Press.

[6] Sigmund, O., 2007. “Morphology-based topology optimization”. Structural and Multidisciplinary Optimization, 33, pp. 401–414.

[7] Dede, E. M., 2009. “Automated design of thermal management structures using topology optimization”. Journal of Mechanical Design, 131(6).

[8] Ignuta-Ciuncanu, M. C., St¨ark, H., and Martinez-Botas, R. F., 2025. “Evolutionary design of conductive pathways using a generative autoencoder”. International Communications in Heat and Mass Transfer, 166, 8, p. 109098.
[9] Ignuta-Ciuncanu, M. C., and Martinez-Botas, R. F., 2026. “Evolutionary design of radial fins for forced and natural convection using generative methods”. International Communications in Heat and Mass Transfer, 171, 2, p. 110027.
[10] Femmer, T., Kuehne, A. J., and Wessling, M., 2015. “Estimation of the structure–performance relationship of 3D printed gyroid structures”. Chemical Engineering Journal, 273, pp. 438–445.
[11] Pietropaoli, M., Ahlfeld, R., Montomoli, F., Ciani, A., and D’Ercole, M., 2017. “Design for additive manufacturing: Internal channel optimization”. Journal of Engineering for Gas Turbines and Power, 139(10), 04, p. 102101.
[12] Pulin, P., et al., 2024. “Optimisation of Heat Exchanger Performance Using Modified Gyroid-Based TPMS Structures”. Processes, 12(12).
[13] Mahmoud, D., Tandel, S. R. S., Yakout, M., Elbestawi, M., Mattiello, F., Paradiso, S., Ching, C., Zaher, M., and Abdelnabi, M., 2023. “Enhancement of heat exchanger performance using additive manufacturing of gyroid lattice structures”. The International Journal of Advanced Manufacturing Technology 2023 126:9, 126, 4, pp. 4021–4036.
[14] Daifalla, E., Shahpar, S., Tristanto, I., and Carta, M., 2024. “Multidisciplinary optimization of gyroid topologies for a cold plate heat exchanger design”. J. of Engineering for Gas Turbines and Power, 146, 12.
[15] Bejan, A., 1978. “General criterion for rating heat-exchanger performance”. International Journal of Heat and Mass Transfer, 21, 5, pp. 655–658.
[16] Almerbati, A., Lorente, S., and Bejan, A., 2018. “The evolutionary design of cooling a plate with one stream”. International Journal of Heat and Mass Transfer, 116, 1, pp. 9–15.
[17] Bejan, A., Lorente, S., and Kang, D. H., 2014. “Constructal design of thermoelectric power packages”. International Journal of Heat and Mass Transfer, 79, 12, pp. 291–299.
[18] Gungor, S., Cetkin, E., and Lorente, S., 2022. “Canopy-to-canopy liquid cooling for the thermal management of lithium-ion batteries, a constructal approach”. International Journal of Heat and Mass Transfer, 182, 1, p. 121918.
[19] Caceres-Gonzalez, R., Dıaz, A. J., Pineda, B., and Sarmiento-Laurel, C., 2025. “Enhancing microchannel thermal performance with fractal design and entropy analysis”. Applied Thermal Engineering, 263, p. 125361.
[20] Ignuta-Ciuncanu, M. C., and Martinez-Botas, R. F., 2025. “Discrete svelteness: Evaluating flow structures in generative constructal design”. BioSystems, 4, p. 105459.
[21] Pietropaoli, M., Montomoli, F., and Gaymann, A., 2018. “Three-dimensional fluid topology optimization for heat transfer”. Structural and Multidisciplinary Optimization. 2018 59:3, 59, 10, pp. 801–812.
[22] Ignuta-Ciuncanu, M. C., Tabor, P., and Martinez-Botas, R. F., 2025. “Generative constructal design of a multiphysics heat sink for managing transient thermal loads”. ASME Journal of Heat and Mass Transfer, 148(2), 11, p. 023301.
[23] Langtangen, H. P., and Logg, A., 2017. Solving PDEs in Python: The FEniCS Tutorial I, 1st ed. Springer Publishing Company, Incorporated.
[24] Ignuta-Ciuncanu, M., Tabor, P., and Martinez-Botas, R., 2024. “A generative design framework for passive thermal control with macroscopic metamaterials”. Thermal Science and Engineering Progress, 51, p. 102637.
[25] Nield, D., and Bejan, A., 2013. Convection in Porous Media, Fourth Edition. Springer, Switzerland.
[26] Guillermo, O. R. L., Arturo, G. O., James, P. B., and Saul, P., 2025. “Computational analysis and engineering modeling for the heat transfer and fluid flow through the gyroid TPMS structure”. Applied Thermal Engineering, 268, 6, p. 125865.
[27] Banthiya, A., Navaresse, B., Pan, L., and Weibel, J. A., 2025. “Simultaneous topology optimization of two hydraulically interconnected porous flow layers in cold plates”. International Journal of Heat and Mass Transfer, 241, p. 126671.
[28] Bejan, A., and Lorente, S., 2008. Design with Constructal Theory. John Wiley & Sons.
[29] ANSYS, Inc., 2024. ANSYS Fluent Theory Guide and User’s Guide. ANSYS Release 2024 R1.
[30] Bejan, A., 2015. “Constructal law: Optimization as design evolution”. Journal of Heat Transfer, 137(6), 06, p. 061003.