A machine-learning based approach for multi-scale optimisation of heat exchangers with lattice-like topology

International Journal of Heat and Mass Transfer 256 (2026) 128035

Contents lists available at ScienceDirect

International Journal of Heat and Mass Transfer

journal homepage: www.elsevier.com/locate/hmt

A machine-learning based approach for multi-scale optimisation of heat exchangers with lattice-like topology

A. Chiodi [a,c,*], A. Alaia [b], E. Lombardi [b], M. Cisternino [b], K. Gkaragkounis [b], A. Ferrero [a], S. Shahpar [c]

[a] Politecnico di Torino, Department of Mechanical and Aerospace Engineering (DIMEAS), Turin, Italy
[b] Optimad S.r.l., Turin, Italy
[c] Rolls-Royce Plc, Department of Fluid Mechanics, Derby, United Kingdom

ARTICLE INFO

Keywords:
Heat exchangers
Generative design
Multi-scale optimisation
Multi-fidelity optimisation
Gyroids
Computational fluid dynamics
Cold plate

ABSTRACT

In this study, we present a novel approach for the multi-scale simulation and optimisation of *Micro Channel Heat Exchangers* (MCHX) with lattice-like topology. By treating the *Triply Periodic Minimal Surface* (TPMS) lattice as an equivalent porous medium, the model enables fast and accurate simulations of full-scale 3D MCHX for industrial applications without requiring high-resolution meshing or computationally demanding high-fidelity CFD runs. Micro- and meso-scale effects are accounted for thanks to variable permeability, Forchheimer, and heat transfer coefficients, modelled as non-linear functions of local flow conditions and lattice geometry. These closure relationships are inferred using a *multi-fidelity* machine learning model, trained on a combination of low- and high-fidelity CFD data. This allows the model to capture the effects of fluid flow phenomena occurring at the smallest scale (such as boundary effects and head pressure losses) without the need to rely on high-fidelity simulations. The presented framework offers a favourable balance between accuracy and cost, enabling optimisation within realistic industrial timelines. As a demonstration, the proposed methodology is applied to the optimisation of a heat exchanger used by Rolls-Royce Plc for the thermal management of high-power electronics in aeronautical applications. In particular, three representative configurations are extracted from the Pareto front, respectively optimised for maximum heat transfer, minimum pressure drop, and a balanced trade-off, thus demonstrating the flexibility of the proposed method in targeting different design priorities.

1. Introduction

Micro Channel Heat Exchangers (MCHX) are *Heat Exchangers* (HX) in which the coolant flows in lateral confinements with the dimension of millimeters. Due to their high volumetric heat flux, compactness, and efficient flow distribution, they offer superior specific performance and reduced weight compared to conventional designs [1–3]. However, the traditional design of MCHXs involves costly, iterative *Research and Development* (R&D) cycles with limited testing budgets, often leading to sub-optimal results. *Multi-Scale Optimisation* (MSO) enables advanced geometries that are unreachable via conventional design; on the other hand, these designs are expensive or impossible to manufacture by means of standard *Computer Numerical Control* (CNC) or injection molding. Meanwhile, the *Technology Readiness Level* (TRL) of *Additive Manufacturing* (AM) techniques such as laser-sintering has rapidly matured, enabling the fabrication of complex geometries with shorter lead times and improved design freedom. While AM unlocks MSO potential for MCHX, it also introduces challenges, such as manufacturability limits (e.g., minimum channel wall thicknesses), thermal distortion, difficult post-processing, and altered surface roughness, affecting fluid flow and heat transfer.

In this work, we present an efficient multi-scale simulation and optimisation strategy for MCHXs based on lattice structures. Though applicable to other families of lattices, we focus on *Triply Periodic Minimal Surfaces* (TPMS), specifically gyroids, with variable geometric parameters across the design space. TPMS offer smooth curvature and high surface-to-volume ratio, thus promoting balanced designs and improved perfusion [4]. However, optimal spatial distribution of parameters such as frequency, wall thickness, and orientation remains challenging due to their strong impact on thermo-hydraulic performance [5,6].

Accurate *Conjugate Heat Transfer* (CHT) simulations of MCHXs require high-resolution models to solve the physics at the micro-scale level. Mesh generation is a major bottleneck: geometric complexity often requires manual clean-up, and resulting grids may suffer from

* Corresponding author at: Politecnico di Torino, Department of Mechanical and Aerospace Engineering (DIMEAS), Turin, Italy.
E-mail address: alessandro.chiodi@polito.it (A. Chiodi).

https://doi.org/10.1016/j.ijheatmasstransfer.2025.128035
Received 3 July 2025; Received in revised form 14 October 2025; Accepted 23 October 2025

Optimisation Nomenclature			
N	population size [–]	m	# of objective functions [–]
J	generic objective function [–]	$\mathbf{W}$	RBF weights matrix [–]
N_C	# of porous domain cell centers [–]	N_{RBF}	# of RBF control nodes [–]
$\{\mathbf{x}_i\}_{i=1}^{N_C}$	porous domain cell centers [m]	$\{\mathbf{x}_j\}_{j=1}^{N_{RBF}}$	RBF nodes [m]
$\boldsymbol{\tau}$	RBF nodes wall thicknesses [m]	$\mathbf{R}$	residuals of the CHT problem [–]
$\mathbf{y}$	CHT discrete state unknowns [–]	$\boldsymbol{\psi}$	porous medium parameters set [–]
Physical Modelling Nomenclature			
T	temperature [K]	ρ	density [kg/m^3]
th	gyroid wall thickness [m]	$\mathbf{th}$	gyroid wall thickness vector [m]
p	pressure [Pa]	p°	total pressure [Pa]
$\overline{p}_{in}$	inlet-averaged pressure [Pa]	$\overline{p}_{out}$	outlet-averaged pressure [Pa]
Δp	unit cell pressure drop [Pa]	$\overline{\Delta p}$	inlet-outlet pressure drop [Pa]
$T_{(s)}^{max}$	max solid phase temperature [K]	$T_{(f)}^{max}$	max fluid phase temperature [K]
$\nabla T_{(s)}^{max}$	max solid temperature gradient[K]	$T_{(f)}^{out}$	fluid phase outlet temperature [K]
$\mathbf{u}$	velocity [m/s]	$\langle\mathbf{u}\rangle^{max}$	max Darcy velocity [m/s]
i	enthalpy [J]	h	heat transfer coefficient [W/m^2 K]
$\dot{Q}$	thermal power [W]	q_w	wall heat flux [W/m^2]
ν	kinematic viscosity [m^2/s]	μ	dynamic viscosity [Pa s]
c	specific heat capacity [J/kg K]	k	thermal conductivity [W/m K]
ϵ	porosity [–]	d	hydraulic diameter [m]
L	reference length [m]	a	specific surface area [1/m]
Ω	reference elementary volume [m^3]	A	wet surface area [m^2]
K	permeability [m^2]	C_F	Forchheimer coefficient [–]
$\boldsymbol{\tau}$	deviatoric stress tensor [Pa]	$\mathbf{S_m}$	porous source [Pa/m]
$\mathbf{D}$	impedance tensor [1/m^2]	$\mathbf{F}$	resistance tensor [1/m]
$\mathbf{x}$	position vector [m]	t	time [s]
f	dimensionless spatial frequency [–]	φ	phase-shift [rad]
Re	Reynolds number [–]	Pr	Prandtl number [–]
C_1	1st exponential fitting constant [–]	C_2	2nd exponential fitting constant [–]
Gaussian Process Nomenclature			
D	dataset	n	# of observations
d	# of latent variables	σ^2	variance
l	length-scale	k_{Mat32}	Matèrn 3/2 Kernel
$m(\mathbf{x})$	mean function	$k(\mathbf{x},\ \mathbf{x}')$	covariance function
Brackets and Subscripts			
$\langle\cdot\rangle$	phase volume averaging	$\|\cdot\|$	Euclidean norm
f	fluid	s	solid
sf	solid-fluid interface	α	generic phase

cell skewness or poor orthogonality, thus affecting stability and convergence. Mesh sizes can range from tens to hundreds of millions of cells. While simulating a single configuration is feasible using *High Performance Computing* (HPC), optimisation runs involve the evaluation of hundreds of variants, making it impractical within industrial timelines.

In recent years, several studies have adopted computationally demanding strategies for the simulation and optimisation of MCHXs. Most rely exclusively on *High-Fidelity* (HF) CFD to evaluate each design. While accurate, this involves enormous computational effort, especially for parametric sweeps or optimisation loops. For instance, [7] conducted detailed CHT simulations of cross-flow MCHXs, capturing fine-scale interactions at the cost of heavy meshing and long runtimes. In [8], HF CFD was used to enhance thermal performance in heat sinks via flexible vortex generators, again requiring intensive computation. In [9], the combination of a response surface method with CFD has been used to optimise air–water MCHXs, showing both accuracy and high computational cost. Similarly, [10] performed multi-objective optimisation of cold plates (based on gyroids) using HF CFD, accounting for pressure drop and heat transfer, but at a significant computational effort. These examples highlight the need for reduced-order modelling as a credible alternative for scaling up the exploration of the design space.

To address the problem of mesh generation, several works published in recent years employ mesh-less methods. For instance, in [11], a thermal-hydraulic characterisation of TPMS periodic structures is presented using a mesh-less *Radial Basis Function - Finite Difference* (RBF-FD) method. The proposed approach results in a high-accuracy assessment of fluid dynamics and heat transfer properties of simple lattices (or single lattice cells) with variable wall-thickness, and is validated against literature results. While the use of mesh-less methods allows bypassing the problem of computing the mesh connectivity, the number of degrees of freedom might still be large since a *dense* point cloud is needed to resolve the smallest space scale. As a consequence, when scaling the mesh-less approach to 3D full-scale HX, the number of problem unknowns increases to roughly the same order of magnitude of standard *Finite Volume* (FV) solvers, rendering optimisations runs still challenging from the computational point of view. Furthermore, in [11] (and other works), only structures with periodic *Boundary Conditions* (BCs) are considered. In real-world problems, boundary effects (i.e., near solid walls) and head pressure losses are not negligible. As a result, the actual lattice permeability may differ significantly from the one computed by means of periodic homogenisation (as demonstrated in Section 5).

In [12], a multi-objective topology optimisation of heat sinks with TPMS lattice is presented based on density mapping approach and validated against experimental results. While the work proposes an elegant and mathematically well-grounded framework, the optimisation is performed for 2.5D lattices, i.e., the optimisation is performed on a 2D domain, and the lattice with variable shape is obtained by extrusion along the 3rd dimension. In our approach, we consider 3D lattices with

variable shape in all the space directions. Furthermore, the approach proposed in [12] still relies on HF simulations, which can be prohibitive expensive for industrial configurations.

Lastly, in [13], a mesh-less method is used to compute the optimal design of a two-fluid TPMS HX. In this work, the frequency and wall thickness of the lattice are encoded by implicit functions, which are the results of the optimisation. This method does not require to generate the mesh, but again relies on HF simulations, which can have prohibitive costs, as discussed earlier.

The innovative aspect of our methodology is two-fold. On one hand, we rely *only* on *Low-Fidelity* (LF) simulations during the optimisation loop. On the other hand, smaller-scale effects (which would require HF computations) are incorporated by means of a multi-fidelity model, which provides an estimate of the equivalent permeability and heat transfer coefficients based on the local average flow and the local shape of the lattice. Our computational meshes are in the order of 10K-100K cells, which is 3 to 4 order of magnitude lower than the typical computational meshes used in HF computations. Only a handful of HF simulations are used to train the multi-fidelity model. With this approximated model, we can obtain predictions with an accuracy similar to HF models, but with a computational cost that is 3 to 4 order of magnitude smaller.

This is achieved by applying homogenisation theory: the lattice matrix is treated as a porous meta-material characterised by permeability and heat transfer coefficients, avoiding direct simulation of the full micro-scale geometry. This philosophy has been used in other fields with satisfactory results. For example, in structural analysis of lattice materials, elasticity equations are solved using "equivalent" Young's modulus and Poisson's ratio [14–16]. In fluid dynamics, creeping flows through porous media are described macroscopically by the Darcy-Forchheimer model [17–19], where both permeability and Forchheimer coefficients depend on geometry and local flow, and are usually determined through HF CFD. A key difference in our case is that local flow regimes span from laminar ($Re \approx 100$) to transitional ($Re \approx 1000$), making inertial effects significant and invalidating classical Darcy–Forchheimer assumptions [20]. To address this, we assume both permeability and heat transfer coefficients to be non-linear functions of the geometry and flow field, leading to a fully non-linear model that requires suitable closure relationships. These are difficult to derive analytically; even simple flows yield closure laws only under restrictive assumptions on geometry and packing (e.g., spherical pores or staggered rods) [21].

In recent years, *Machine Learning* (ML) has obtained tremendous importance due to the technological advancements in training algorithms and data availability. Many techniques originally developed to solve computer vision problems are now employed as routines in more traditional disciplines, including CFD (see [22,23] for a general review). Among others, two promising applications are the use of ML models to infer closure relationships from *Direct Numerical Simulation* (DNS) data and the acceleration of CFD simulations. Some examples include the determination of transport coefficients in turbulence modelling [24], or acceleration of CFD runs in turbomachinery applications (see [25,26], and [27]).

In this work, we propose a multi-fidelity framework to provide closure relationships for the homogenised model. The method is based on the *Non-Linear Auto-Regressive Gaussian Process* (NARGP) [28,29], linking macroscopic lattice properties to the non-linear micro-scale flow behaviour. The ML model bridges the fluid dynamics at the level of the gyroid unit cells (micro-scale) and the averaged equations on the entire MCHX (macro-scale). It is trained on abundant LF data and a smaller set of HF data for accurate predictions [30,31]. While the multi-fidelity model offers a remarkable speed-up of CHT runs and enables an efficient exploration of the design space, it also has drawbacks. The homogenisation procedure, with the consequent approximation of the complex micro-scale flow with an equivalent porous medium, neglects micro-scale phenomena (e.g., recirculation zones, secondary flows, velocity gradients), introducing a trade-off between fidelity and computational tractability. However, the learning framework is physically grounded and interpretable: unlike black-box models, it reflects the structure of the governing equations. Moreover, the integration of HF data enhances accuracy, especially where LF data fail to capture boundary effects or head pressure losses. These benefits are evident in transitional regimes and non-uniform lattice configurations. Thus, the method balances efficiency, accuracy, and physical consistency, making it well-suited for industrial design loops. To test it, we optimise the TPMS-based lattice of a HX used in the thermal management of aeronautical high-power electronics, a test case provided by Rolls-Royce Plc as proof-of-concept for real-world use.

The structure of the present work is presented hereafter. Section 2 focuses on the overall optimisation workflow, with emphasis on the rationale behind the choice of the optimisation algorithm. Section 3 describes the macro-scale mathematical model, based on the equivalent porous media approximation. Sections 4 and 5 describe the thermal/fluid dynamic characterisation of individual lattice cells. In Sections 6 and 7, we discuss the methodology used to incorporate HF data into the characterisation process. We also explain how LF and HF data are integrated in the multi-fidelity model. The application of the proposed methodology to our target test case is presented in Section 8. Finally, concluding remarks and future perspectives are summarised in Section 9.

2. Optimisation strategy

In a general framework, the aim of optimisation is to compute the best distribution of gyroid lattice parameters, such as wall thickness, local orientation, and anisotropic frequency, to minimise or maximise given objectives under constraints. Lattice anisotropy and orientation can improve fluid perfusion and reduce pressure drop, but may also reduce coolant residence time, adversely affecting heat transfer. Objectives and constraints are typically defined using macro-scale integral quantities. Common goals include maximising thermal exchange and efficiency, and minimising pressure losses. Constraints often enforce upper limits on temperature and its gradient to avoid thermal cracking, and manufacturing constraints are imposed via bounds on wall thickness.

This work focuses on thermo-fluid optimisation, though the same workflow can be extended to structural or process simulations. Assigning lattice parameters per computational cell leads to high-frequency variations and checkerboard patterns [32], which hinder manufacturability. To avoid this, filtering strategies are used. We adopt a reduced-DOF approach: lattice parameters are optimised at a subset of control nodes scattered in the domain. Using *Radial Basis Functions* (RBF), we interpolate these values across all cells. This suppresses checkerboarding and speeds up convergence. Once interpolated, a macro-scale CHT simulation is run. During each CFD iteration, the multi-fidelity ML model is queried to compute local permeability, Forchheimer, and heat transfer coefficients. The results feed the objective and constraint functions, which are passed to the optimiser. A schematic of the workflow is shown in Fig. 1.

The optimisation workflow is implemented in modeFRONTIER (ESTECO, [33]), using the proprietary HYBRID algorithm, which combines a steady-state *Non-Dominated Sorting Genetic Algorithm* (NSGA-II) with *Sequential Quadratic Programming* (SQP) [34]. NSGA-II was chosen since it alleviates the main limitations of earlier multi-objective evolutionary strategies with non-dominated sorting. In particular, previous methods suffered from high computational complexity in ranking solutions ($O(mN^3)$, with m objectives and N population size), the absence of an explicit elitism mechanism to preserve the best individuals, and the need for a problem-dependent tuning of a sharing parameter to control solution diversity along the Pareto front. NSGA-II overcomes all the above-mentioned difficulties [35], by introducing a "fast" non-dominated sorting approach with $O(mN^2)$ complexity, an explicit elitist

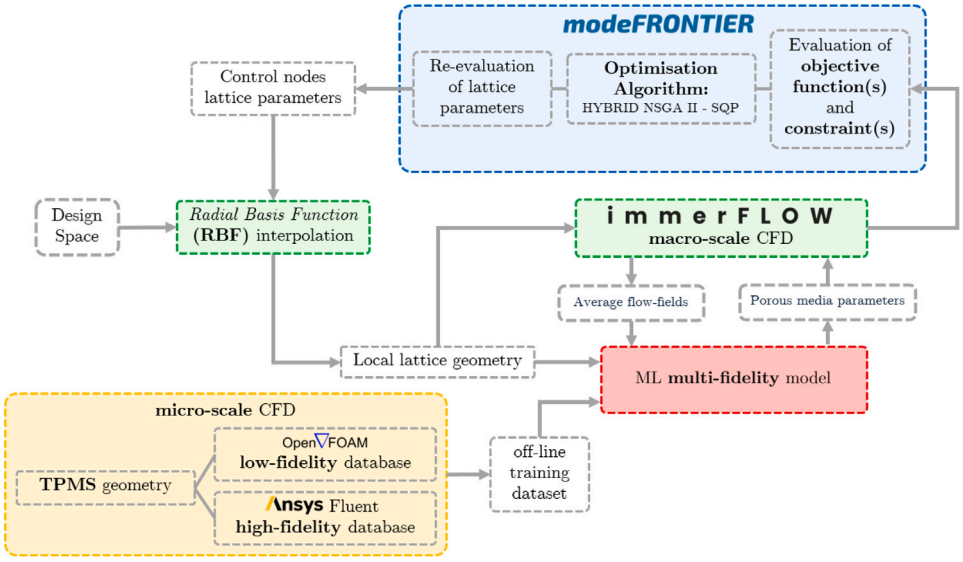

Fig. 1. Schematic depiction of the multi-scale optimisation workflow.

approach, and the crowding-distance metric that preserves diversity without requiring user-defined similarity thresholds. With those devices, both diversity and the spread of the solutions on the Pareto front are guaranteed without the need for problem-specific tuning.

Unlike conventional hybrid approaches, where SQP is employed only after the genetic algorithm has converged, the HYBRID framework integrates the *Adaptive Filter SQP* (AFilterSQP) directly within the evolutionary cycle, as a third operator besides crossover and mutation [36]. During optimisation, selected non-dominated individuals from the NSGA-II population undergo local refinement through SQP, which, thanks to its filtering strategy, accepts only iterations that improve objectives or reduce constraint violations. The refined solutions are continuously reinjected into the NSGA-II population, ensuring a dynamic exchange of information: the genetic algorithm provides diversity and broad exploration, while SQP contributes locally converged, high-quality solutions. This parallel cooperation enhances convergence compared with the traditional sequential "GA plus refinement" approach.

The motivations behind the choice of the HYBRID optimisation algorithm (NSGA-II + SQP) over a traditional gradient-based method lie in its low intrusiveness, its ability to rapidly explore the parameter space and detect global minima, and its robustness when dealing with complex, non-differentiable objective functions or constraints. In addition, HYBRID offers high scalability and is naturally suited for problems involving multiple objectives in multi-disciplinary settings. By contrast, gradient-based optimisation typically requires ad hoc regularisation strategies to handle non-differentiable objectives and often suffers from convergence difficulties in the presence of noisy gradients.

3. Macro-scale modelling

At the macro-scale, we identify three regions in the HX: a fluid region which contains only the fluid phase (e.g., inlet/outlet port), a solid region which contains only the solid phase (e.g., the outer case of the HX), and the porous matrix, which contains both the solid and fluid phases. In the solid phase, we integrate only the energy equation for a static phase, whereas in the fluid region, we integrate the incompressible *Navier–Stokes* equations, including the energy equation. Generally speaking, while the Reynolds number within the gyroid lattice can vary between laminar to transitional, fully turbulent flow can sometimes be observed in the inlet/outlet channels of the MCHX. In the present study, we only focus on transitional flow, i.e., the *Boundary Conditions* (BCs) are such that a fully turbulent flow never develops in the inlet/outlet channels.

The porous matrix represents the porous lattice whose parameters are subjected to optimisation. In this region, the model equations are the ones stemming from the standard theory of multi-phase flows through porous media, and they are obtained by suitable volumetric averages of conservation laws for mass, momentum, and energy over a *Reference Elementary Volume* (REV) [17]. In our case, the REV Ω is represented by a single lattice cell. It consists of two separate phases, solid ($\Omega_{(s)}$, with boundary $\partial\Omega_{(s)}$), and fluid ($\Omega_{(f)}$, with boundary $\partial\Omega_{(f)}$), s.t. $\overline{\Omega} = \overline{\Omega_{(s)}} \cup \overline{\Omega_{(f)}}$, where the bar indicates the closure of the set. From now on, subscripts (f) and (s) will be used to indicate fluid and solid phases, respectively. It is possible to introduce the *phase volume-averaging* [17] (indicated as $\langle - \rangle$) of an arbitrary field quantity $\phi_{(\alpha)}(\mathbf{x}, t)$, related to the α phase, as follows:

$$\left\langle \phi_{(\alpha)} \right\rangle (t) = \frac{1}{\Omega} \int_{\Omega_{(\alpha)}} \phi_{(\alpha)}(\mathbf{x}, t)\, d\Omega_{(\alpha)} \tag{1}$$

The vast majority of studies on heat and fluid flow through porous media rely on Darcy's law, which relates the *phase volume-averaged* fluid velocity $\langle \mathbf{u} \rangle$ (called Darcy velocity) to the pressure gradient $\nabla \langle p \rangle$ as follows [17]:

$$\langle \mathbf{u} \rangle = -\frac{K}{\mu_{(f)}} \nabla \langle p \rangle \tag{2}$$

where, K is the permeability tensor and $\mu_{(f)}$ the dynamic viscosity of the fluid. Since the solid phase is stationary, we will denote fluid-phase averages simply as $\langle \mathbf{u} \rangle$ and $\langle p \rangle$ throughout this work. This principle remains the foundation of most porous media flow models, where permeability K serves as an analogue to thermal conductivity in heat transfer. Deviations from Darcy's law occur when the local Reynolds number Re, based on the mean pore diameter d, exceeds unity [17]. The local Reynolds number and mean pore diameter are defined as follows:

$$Re = \frac{\rho_{(f)} \| \langle \mathbf{u} \rangle \| d}{\mu_{(f)}}, \quad d = \frac{4\epsilon}{a_{(s,f)}} \tag{3}$$

where ϵ is the porosity, and $a_{(s,f)} = A/\Omega$ is the specific surface area (related to the wet surface, A). In this regime, inertial effects become significant as boundary layers develop along solid surfaces, leading to a non-linear relationship between the pressure gradient and velocity. To account for these inertial effects, Darcy's model is modified by introducing the well-known Darcy–Forchheimer correction [17]:

$$-\nabla \langle p \rangle = \frac{\mu_{(f)}}{K} \langle \mathbf{u} \rangle + \rho_{(f)} \frac{C_F}{\sqrt{K}} \| \langle \mathbf{u} \rangle \| \langle \mathbf{u} \rangle \tag{4}$$

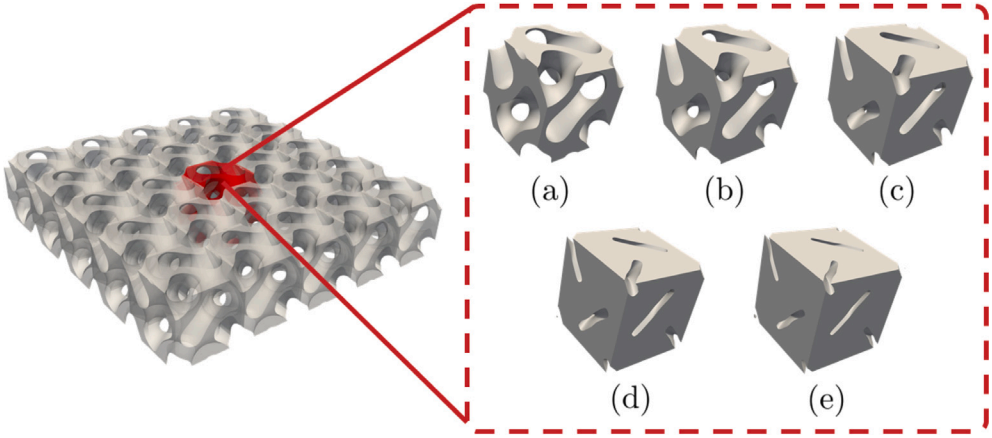

Fig. 2. Schematic depiction of the periodic TPMS lattice, together with the gyroid unit cell at different wall thicknesses. (a) $th = 1.25\,\text{mm}$, (b) $th = 2\,\text{mm}$, (c) $th = 3\,\text{mm}$, (d) $th = 3.5\,\text{mm}$, and (e) $th = 3.75\,\text{mm}$.

where $\rho_{(f)}$ denotes the fluid density. On the right-hand side of Eq. (4), the first term represents the viscous (Darcy) contribution, while the second term accounts for inertial effects through the Forchheimer coefficient C_F. At the scale of the HX, besides the incompressibility condition, the fluid flow is governed by , where $\mathbf{S_m}$ acts as a sink term:

$$\begin{cases} \nabla \cdot \langle \mathbf{u} \rangle = 0 \\ \dfrac{\partial \langle \mathbf{u} \rangle}{\partial t} + (\langle \mathbf{u} \rangle \cdot \nabla) \langle \mathbf{u} \rangle = \dfrac{1}{\rho_{(f)}} \left(-\nabla p + \mu_{(f)} \nabla^2 \langle \mathbf{u} \rangle + \mathbf{S_m} \right) \end{cases} \tag{5}$$

In the previous equation $\mathbf{S_m}$ is the Darcy–Forchheimer momentum sink term, expressed as:

$$\mathbf{S_m} = \left(\mu_{(f)} \mathbf{D} \langle \mathbf{u} \rangle + \frac{1}{2} \rho_{(f)} \, \text{tr}(\langle \mathbf{u} \rangle \cdot \mathbf{I}) \, \mathbf{F} \right) \langle \mathbf{u} \rangle \tag{6}$$

which is the generalised tensor form of Eq. (4) where $\mathbf{D}$ is the impedance tensor, and $\mathbf{F}$ is the resistance tensor (also referred to as Forchheimer tensor [37]). For the isotropic case, $\mathbf{D}$ and $\mathbf{F}$ can be simplified as $D_{i,i} = K_{i,i}^{-1}$, and $F_{i,i} = 2\, C_{Fi,i} / \sqrt{K_{i,i}}$.

The energy equations are formulated separately for the fluid and solid phases. These equations characterise the evolution of phase volume-averaged temperature fields ($\langle T_{(f)} \rangle$ and $\langle T_{(s)} \rangle$) within each phase, capturing the interfacial convective heat transfer. Mathematically, they are expressed as follows [38]:

$$\begin{cases} \epsilon \rho_{(f)} c_{(f)} \left(\dfrac{\partial \langle T_{(f)} \rangle}{\partial t} + \langle \mathbf{u} \rangle \cdot \nabla \langle T_{(f)} \rangle \right) = \nabla \cdot \left(k_{(f)} \nabla \langle T_{(f)} \rangle \right) + \\ \qquad\qquad + \, h_{(s,f)} a_{(s,f)} \left(\langle T_{(s)} \rangle - \langle T_{(f)} \rangle \right) \\ (1-\epsilon) \rho_{(s)} c_{(s)} \dfrac{\partial \langle T_{(s)} \rangle}{\partial t} = \nabla \cdot \left(k_{(s)} \nabla \langle T_{(s)} \rangle \right) \quad - h_{(s,f)} a_{(s,f)} \left(\langle T_{(s)} \rangle - \langle T_{(f)} \rangle \right) \end{cases} \tag{7}$$

Here, the coefficients $c_{(f)}$, $c_{(s)}$, $k_{(f)}$, and $k_{(s)}$ are the specific heat capacities and thermal conductivities for the fluid and solid phases, respectively. $h_{(s,f)}$ is the heat transfer coefficient. In standard porous media, low Reynolds numbers Re allow inertial effects to be captured via the Forchheimer term. In contrast, local Re in lattice channels can exceed 1000, requiring K, C_F, and $h_{(s,f)}$ to be treated as non-linear functions of local geometry and Reynolds number.

are numerically integrated at the macro-scale in ImmerFLOW (Optimad), a *Finite Volume* (FV) CFD solver based on the immersed boundary conditions paradigm to efficiently handle complex geometries [39]. To avoid the cost of a monolithic solver, the temperature equations of the two phases are weakly coupled. An initial guess for the second-phase temperature provides the source term for the first. Each equation is solved using pseudo-time stepping until the steady state is reached, and the updated solution is used to refresh the other phase. This loop continues until both pseudo-time derivatives converge. For the momentum , which includes convective, Darcy, and Forchheimer terms, fixed-point iterations are used: starting from a velocity guess, the ML model provides K and C_F; these are updated at each step until convergence.

A final remark deals with the validity of the Darcy–Forchheimer model, which requires strictly positive K values. Permeability equal to zero implies blocked channels, while negative values are unphysical. While pseudo-density optimisation methods [32] often apply near-zero values, these may introduce numerical stiffness requiring the usage of implicit solvers. Instead, we impose lower bounds on the channels' wall thickness during optimisation, thus avoiding blocked paths, preserving solver compatibility, and meeting manufacturing limits, though this restricts the formation of fully solid barriers in the porous matrix, like septa or dividers.

4. Micro-scale modelling

To train the ML model for the prediction of porous media parameters (i.e., the entries of $\mathbf{D}$, $\mathbf{F}$, and $h_{(s,f)}$), a database of CFD simulations at the micro-scale was created for lattice cells with varying thicknesses and flow conditions. The present work focuses on gyroids, a subset of TPMS, which are particularly suitable for CFD applications, especially CHT, due to their smooth curvature and high surface-to-volume ratio, which enhances the convective heat transfer. In transitional/turbulent regimes, the gyroid's internal shape promotes mixing via turbulent advection, increasing both pressure drop and heat transfer. TPMS geometries are also structurally advantageous: the absence of sharp corners prevents stress concentrations. The lattice geometry was generated using MIMIC (Optimad) [40]. The equation of the gyroid is:

$$\begin{aligned} f(x,y,z) = & \sin\left(\frac{2\pi f_x}{L} x + \varphi_x\right) \cos\left(\frac{2\pi f_y}{L} y + \varphi_y\right) \\ & + \sin\left(\frac{2\pi f_y}{L} y + \varphi_y\right) \cos\left(\frac{2\pi f_z}{L} z + \varphi_z\right) \\ & + \sin\left(\frac{2\pi f_z}{L} z + \varphi_z\right) \cos\left(\frac{2\pi f_x}{L} x + \varphi_x\right) \end{aligned} \tag{8}$$

where f_x, f_y, and f_z are the spatial frequencies along the coordinate axes, while φ_x, φ_y, and φ_z represent the phase shifts, and $L = 1\,\text{cm}$ is the reference length of the unit cell. Gyroid with varying wall thickness can be generated by computing the *Signed Distance Function* (SDF) of the 0 - contour of Eq. (8) and then extracting the iso-contour corresponding to the desired thickness from the SDF (see Fig. 2).

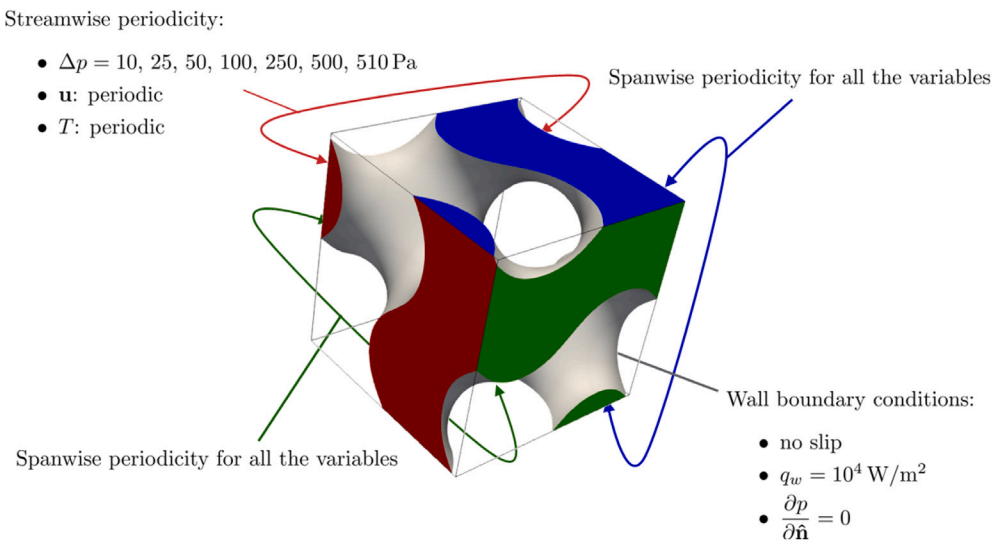

Fig. 3. Summary of the *Boundary Conditions* (BCs) on each gyroid unit cell.

To estimate the equivalent permeability, Forchheimer, and heat transfer coefficient for each gyroid geometry, two types of CFD simulations were performed. First, a cold run is conducted, in which only the continuity and momentum equations are solved to characterise the impedance and resistance tensors, **D** and **F**, as a function of the pressure gradient acting across the lattice cell. Then, a hot run is performed to characterise the heat transfer coefficient $h_{(s,f)}$, starting from the steady-state solution of the cold run. In this second simulation, the temperature field is introduced by imposing a fixed heat flux BC at the gyroid walls, while the momentum solution is kept unchanged. Since the thermo-physical properties of the coolant are assumed to be constant with respect to the temperature, the momentum and energy equations are completely decoupled. As a result, the heat transfer coefficient can be considered approximately constant with respect to temperature variations.

Both hot and cold runs were performed using the `buoyantPimpleFoam` solver from the OpenFOAM suite [41], where the implemented model equations are the standard *Navier–Stokes* equations for unsteady, laminar, and incompressible flows. The energy equation is written in terms of the enthalpy i.

$$\begin{cases} \nabla \cdot \mathbf{u} = 0 \\ \dfrac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\,\mathbf{u} = -\dfrac{1}{\rho_{(f)}} \nabla p + \nu_{(f)} \nabla^2 \mathbf{u} \\ \rho_{(f)} \left[\dfrac{\partial i}{\partial t} + \mathbf{u} \cdot \nabla i + \dfrac{\partial}{\partial t} \left(\dfrac{1}{2} \|\mathbf{u}\|^2 \right) + \nabla \cdot \left(\dfrac{1}{2} \|\mathbf{u}\|^2 \mathbf{u} \right) \right] - \dfrac{\partial p}{\partial t} = \dfrac{k_{(f)}}{c_{(f)}} \nabla^2 i \end{cases} \tag{9}$$

In the target application of this work, the fluid phase is water. The following properties were considered in the cold/hot runs: - density $\rho_{(f)} = 998.2\,\mathrm{Kg/m^3}$ - specific heat capacity $c_{(f)} = 4,182\,\mathrm{J/Kg\,K}$ - dynamic viscosity $\mu_{(f)} = 1.003 \times 10^{-3}\,\mathrm{Pa\,s}$ - kinematic viscosity $\nu_{(f)} = 1.0047 \times 10^{-6}\,\mathrm{m^2/s}$ - thermal conductivity $k_{(f)} = 0.6\,\mathrm{W/m\,K}$ - Prandtl number $Pr = 6.9909$. In cold runs, a pressure gradient is imposed along one lattice direction at a time to characterise the corresponding row of the permeability tensor. Periodic BCs on pressure are applied in the remaining two directions orthogonal to the gradient. Periodic BCs on velocity are used in all directions. This setup corresponds to OpenFOAM `cyclicAMI` BCs and is commonly used to emulate fully developed flows in the bulk region of a porous medium [41]. For low Reynolds numbers ($Re = 0 - 300$), the Darcy velocity $\langle u \rangle$ is computed as the (phase) volumetric average of the velocity component along the pressure gradient at the steady state. (see Fig. 4(a)). Starting from $Re > 400$, numerical simulations at the micro-scale exhibit unsteady behaviour. For these cases, the Darcy velocity was computed as the *phase volume-average* of the velocity component along the pressure gradient averaged over a time window of 1 s at statistical convergence. In the performed numerical experiments, statistical convergence is reached after approximately 1.5 s (see Fig. 4(c)). Hot runs are performed starting from $t = 1.5$ s. The initial temperature assigned to the domain is $T_0 = 293\,\mathrm{K}$, and `cyclicAMI` BCs are imposed on the external patches (i.e., inlet, outlet, top, bottom, front, and back). A fixed heat flux is imposed at the wet surface, A, equal to $q_w = 10^4\,\mathrm{W/m^2}$. A summary of the applied BCs can be found in Fig. 3. The heat transfer coefficient $h_{(sf)}$ is computed as:

$$h_{(sf)}(t) = \frac{\dot{Q}(t)}{A\left(T_A(t) - \left\langle T_{(f)} \right\rangle(t)\right)} \tag{10}$$

where $\dot{Q}$ is the total heat power exchanged on the wet surface A, T_A is the surface-averaged wall temperature, and $\left\langle T_{(f)} \right\rangle$ the phase volume-averaged fluid temperature. After an initial transient due to the impulsive activation of the thermal solver, $h_{(sf)}$ reaches a steady value for $Re = 0 - 300$ (Fig. 4(b)). For higher Re, similarly to Darcy velocity estimation, it is averaged over a 1 s time window once statistical convergence is reached (Fig. 4(d)).

Simulations use a first-order implicit time scheme (`Euler`), second-order gradient (`Gauss linear`), upwind divergence (`Gauss upwind`), and second-order conservative Laplacian (`Gauss linear corrected`) [41]. A total of 120 cases were run on $L = 1$ cm gyroid lattices with 15 wall thicknesses (0.25 – 3.75 mm, with steps of 0.25 mm) and 8 pressure drops (0, 10, 25, 50, 100, 250, 500, and 510 Pa). Each case ran in parallel on Optimad's HPC cluster using 12 MPI processes and 30 GB RAM. Representative results for 0.5 and 2.5 mm wall thicknesses are shown in Fig. 4, while velocity fields appear in Fig. 5. Meshes were generated with `snappyHexMesh`, using 500,000 mixed cells (tetrahedral and hexahedral), and 7 prism layers with a first layer thickness of 0.1 relative to the closest undistorted cell size outside the layer, and expansion ratio of 1.2.

5. Lattice-cells characterisation

The equivalent permeability K and Forchheimer coefficient C_F are computed by inverting the Darcy–Forchheimer law written along the direction of the external pressure drop applied to the gyroids:

$$\frac{\Delta p}{\mu_{(f)} \langle u \rangle L} = -\frac{1}{K} - \frac{C_F}{\sqrt{K}} \frac{\rho_{(f)} \langle u \rangle}{\mu_{(f)}} \tag{11}$$

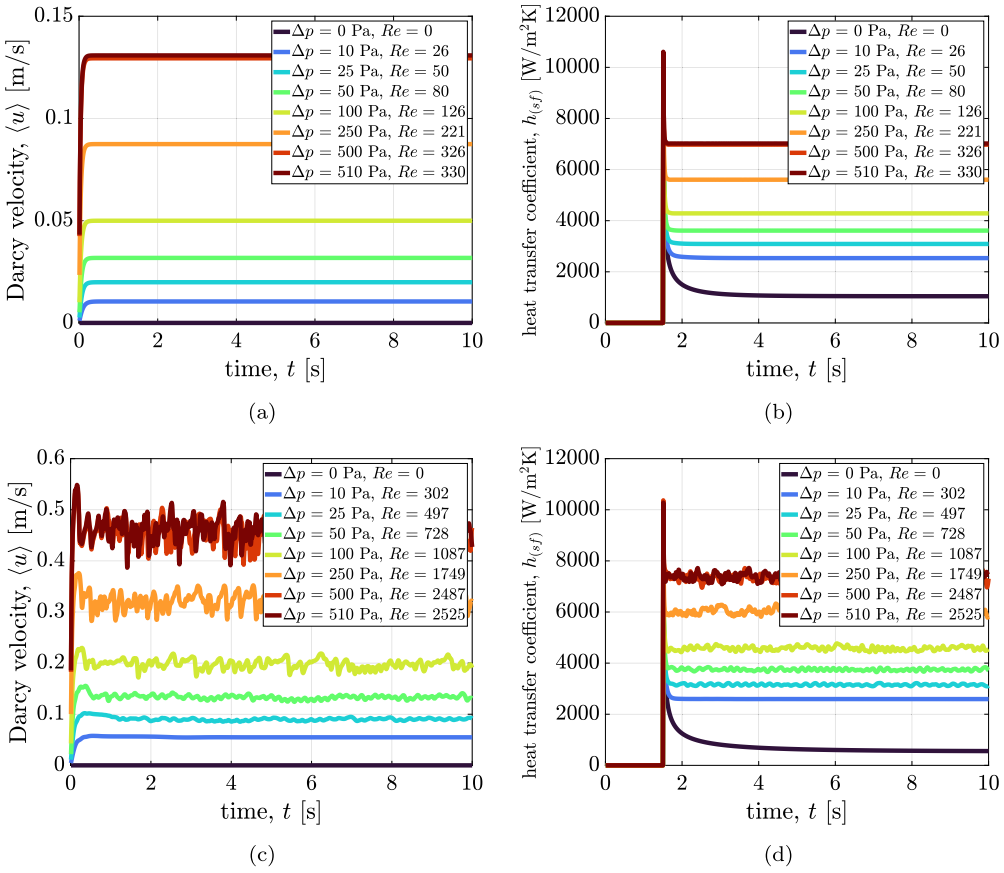

Fig. 4. Results of Darcy velocity, $\langle u \rangle$ (left), and heat transfer coefficient, h_{sf} (right), for a gyroid with wall thickness 2.5 mm (top row) and 0.5 mm (bottom row).

which corresponds to a straight line in the plane of $\Delta p/(\mu_{(f)} \langle u \rangle L)$ versus $\rho_{(f)} \langle u \rangle / \mu_{(f)}$. Therefore, K and C_F can be computed by simply estimating the intercept and slope of the above line. To do that, only a single additional point is required. However, instead of relying on a single auxiliary point, we decided to use 4 additional points to account for noise in the training data. The 4 auxiliary points are selected in a close neighbourhood of a query point, $(\Delta p, \langle u \rangle)$, as shown in Fig. 6(b). Knowing the value of the Darcy velocity $\langle u \rangle$ at such additional points allows us to compute the line that fits the auxiliary points in a least-squares sense. Next, the values of K and C_F can be estimated from the slope and intercept of this line. However, since the auxiliary points are not part of the dataset, the Darcy velocity is first interpolated in the space $(th, \Delta p)$. A preliminary study showed that the trend is well described by the following exponential fitting:

$$\langle u \rangle^* = C_1 \, \Delta p^{*C_2} \tag{12}$$

where $\langle u \rangle^*$ is the resultant Darcy's velocity at the generic pressure drop Δp^*, while $C_1 = C_1(th)$ and $C_2 = C_2(th)$ are two constants depending on the wall thickness th (see Fig. 6(a)).

5.1. Preliminary validation using low-fidelity (LF) data

The characterisation described in the previous section relies solely on LF simulations, i.e., CFD runs performed on individual lattice cells assuming an infinite periodic lattice. However, lattice cells near the solid walls exhibit lower permeability due to boundary effects. As such, we expect the homogenised model to underestimate the pressure drop across the full HX. To verify this, we ran CHT tests on a simplified HX composed of a $4 \times 4 \times 1$ gyroid matrix ($L = 1$ cm, $th = 1.5$ mm). The domain includes an aluminium case with two heat sources. The material properties of aluminium are: - density $\rho_{(s)} = 2719\,\text{kg/m}^3$ - thermal conductivity $k_{(s)} = 202.4\,\text{W/m K}$ - specific heat capacity $c_{(s)} = 871\,\text{J/kg K}$. A pressure gradient drives the flow across the HX. HF simulations used the steady incompressible *Reynolds Averaged Navier–Stokes* (RANS) solver in ANSYS Fluent [42], with the *k–ω Shear Stress Transport* (SST) model [43], suitable for capturing local unsteadiness when $Re > 400$. Total pressure drops tested were $\Delta p° = 100$, 500, and 1000 Pa. The heat equation was solved in the solid domain with a uniform heat flux $q_w = 94{,}500\,\text{W/m}^2$ on the heating plates. The solid-fluid interface was treated as thermally coupled; other walls were adiabatic. Gradient approximation used `Least Squares Cell Based`, `Second Order` for pressure, `Second Order Upwind` for momentum, and `First Order Upwind` for k and ω equations. The mesh, generated in BOXER [44], contains approximately 63 million mixed cells with 8 prismatic layers. The first layer thickness was 10% of the adjacent undistorted cell, with an expansion ratio of 1.1. Simulations ran for 1000 iterations on 80 Intel(R) Xeon(R) Gold 6148 CPUs to ensure convergence.

The homogenised model used the same domain and BCs. However, the lattice was replaced with a porous medium, where permeability and resistance depend on local Darcy velocity and gyroid wall thickness. These coefficients were inferred using an RBF regressor trained on LF data. Due to their non-linear dependence, fixed-point iterations have been used (see Section 3). Impedance and Forchheimer's tensors were computed at 16 control points (centres of gyroid cells), with values elsewhere interpolated via RBFs (see Fig. 7). Mass flow rates from HF simulations were imposed as inlet BCs for the homogenised model, and the resulting pressure drops were compared. Fig. 11 shows that the homogenised model consistently underestimates the pressure losses, meaning that the LF-trained regression overestimates permeability. This underestimation arises from two factors: *Boundary effects*, cells at the periphery are less permeable due to nearby solid walls, which

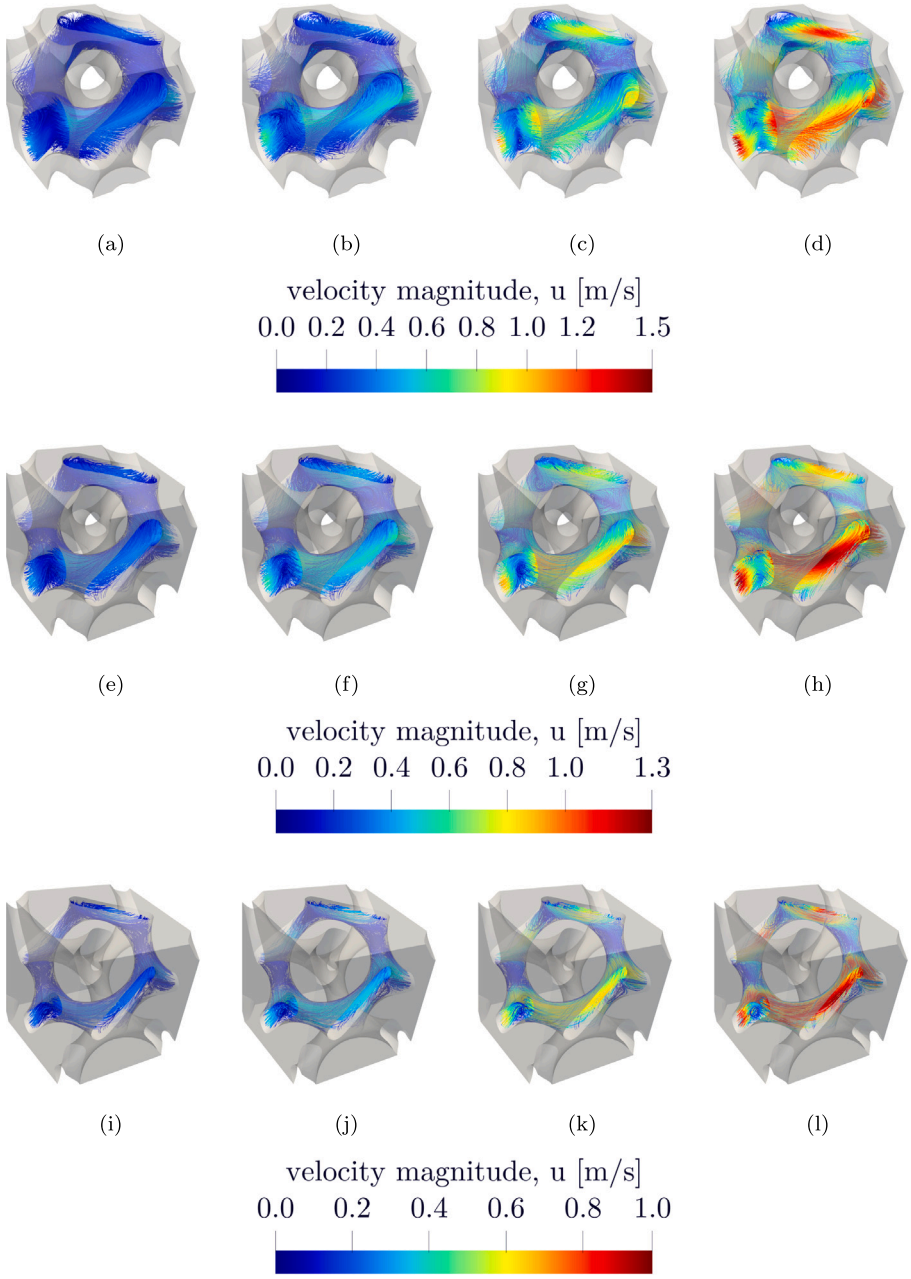

Fig. 5. Velocity magnitude fields at pressure drops $\Delta p = 50, 100, 250, 500$ Pa for a gyroid with wal thickness equal to 1.25 (first row), 2 (second row), and 3 mm (third row).

are ignored in LF runs, where the unit cells are provided by periodic BCs. *Head losses*, at inlet/outlet transitions, aligned channels saturate early, forcing redistribution and causing local pressure rises, which the homogenised model cannot capture due to its uniform interface.

To improve the accuracy, LF data were enriched with HF simulation results on a full HX. Permeability and Forchheimer coefficients were extracted for each cell and combined with LF data to train separate multi-fidelity models for 5 gyroid cell families: boundary, inlet, outlet, bulk, and inlet-wall-adjacent cells. Given the cost of HF runs, only a limited number of gyroid wall thicknesses and flow conditions were assessed. LF data gives the general trends, while HF data is used as a correction to enhance the model fidelity. As with K and C_F, the heat transfer coefficient $h_{(sf)}$ depends on wall thickness and flow conditions. However, since solid walls and inlet/outlet transitions have minimal impact on $h_{(sf)}$, a regression only based on LF data is sufficient to get accurate results (see Section 7). The HF data generation strategy is detailed in the next section, while the multi-fidelity model is discussed in Section 7.

6. Generation of the *high-fidelity* (HF) database

Due to the computational cost of HF runs, only a subset of gyroid wall thicknesses and flow conditions are considered for the building of the HF database. Specifically, we considered the following wall thicknesses: 0.25, 0.75, 1.5, 2.5, 3.25, and 3.75 mm, which represent a good compromise between computational cost and a good coverage of the parameter space. The following total pressure drops are applied across the HX: 10, 25, 50, 100, 500, 1000, and 2000 Pa to guarantee sufficient variability. Simulations are performed in ANSYS Fluent using the same setup described in Section 5. Examples of HF results are shown in Fig. 8 for different wall thicknesses, at $\Delta p^\circ = 2000$ Pa.

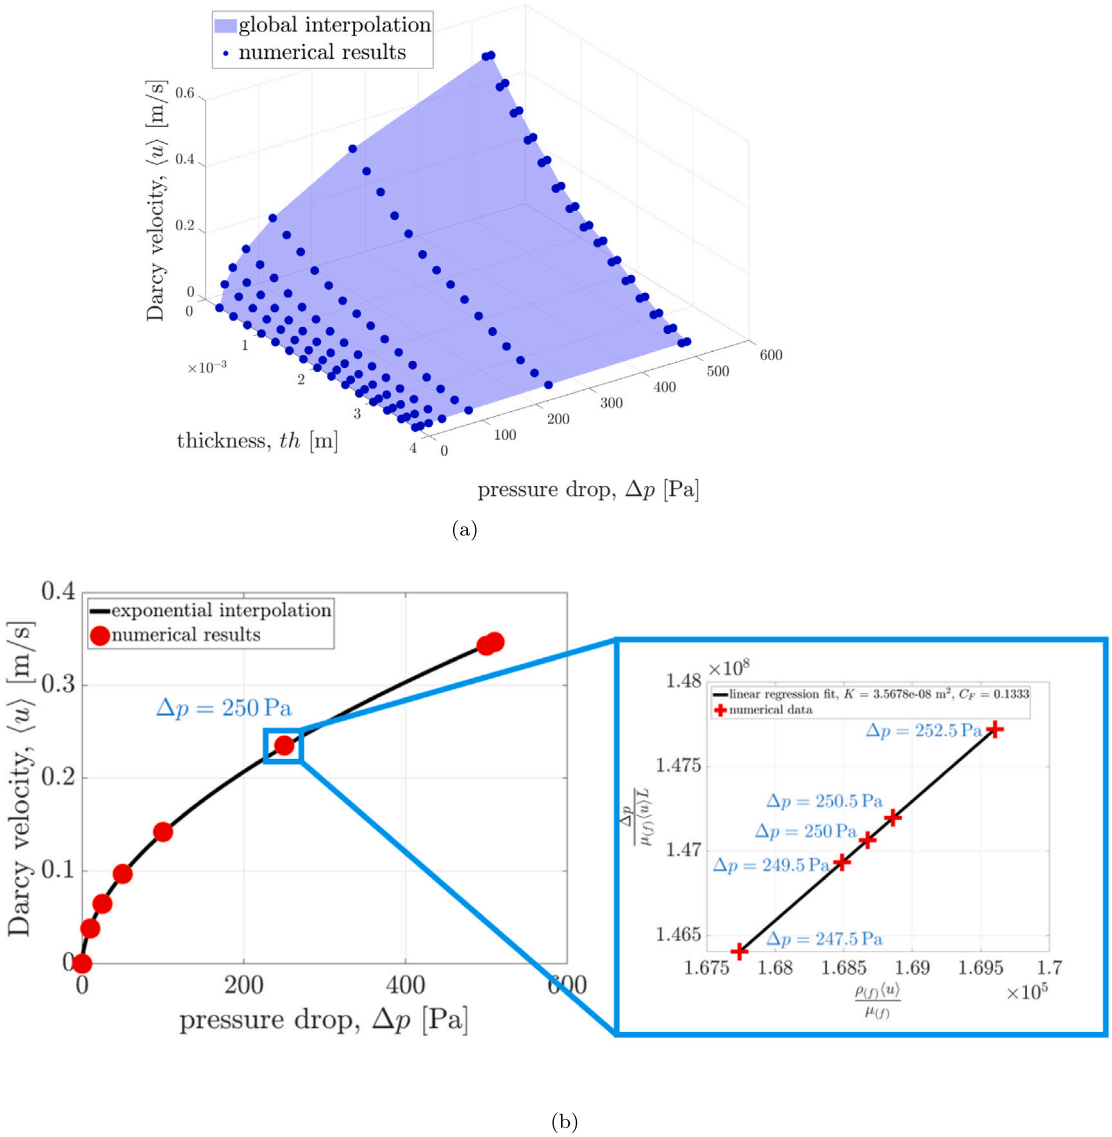

Fig. 6. (a) Exponential fitting of the Darcy velocity $\langle u \rangle$ as a function of Δp in the parameter space $th - \Delta p$. (b) Estimation of K and C_F for a wall thickness 1.5 mm and $\Delta p = 250$ Pa. Given a point $(\Delta p, \langle u \rangle)$ in the dataset, we consider four neighbouring points (i.e., $\Delta p = 247.5$, $\Delta p = 249.5$, $\Delta p = 250.5$, and $\Delta p = 252.5$ Pa). The value of the Darcy velocity at such points is computed using the exponential fitting in Eq. (12). K and C_F are then computed as the slope and intercept of the line which best fits (in a least-square sense) the auxiliary points and the query point.

For each of the 16 gyroid cells in the lattice matrix, the Darcy velocity $\langle u \rangle$ was computed using the same approach described in Section 5. The main difference here is that the characterisation is performed separately for each cell, depending on its location within the lattice. To this end, we distinguish 5 families of lattice cells (see Fig. 9(a)): cells close to solid boundaries; cells located near the inlet; cells located at the outlet; cells located near the inlet *and* close to the solid walls; and cells located within the bulk of the lattice. For each of these families, we build a different multi-fidelity model using LF characterisation data together with K and C_F extracted from the HF simulations. Fig. 9(b) shows that solid walls and head pressure losses are responsible for a reduction of the lattice permeability. The effect of boundary walls and head pressure losses is even more evident by looking at Fig. 10, where K and C_F computed from both LF and HF data are compared. In particular, it is shown that the effect of including solid walls and head pressure losses is to reduce the permeability coefficient and increase the inertial effects, as reflected by the larger Forchheimer's coefficient.

Similarly, the heat transfer coefficient $h_{(sf)}$ is a function of the local lattice geometry (i.e., the gyroid wall thickness) and flow conditions, so, in principle, a multi-fidelity model to regress the heat transfer coefficient from the local flow conditions should be built as well. However, in our numerical experiments, we found that a single-fidelity RBF model is accurate enough (see Table 1). The data obtained from both HF and LF computations are used for the *off-line* training of the multi-fidelity model, which will be discussed in the next section.

7. Multi-fidelity model

To infer closure relationships for lattice coefficients, we adopt a multi-fidelity regression model based on the *Non-Linear Auto-Regressive Gaussian Process* (NARGP) model [29]. This framework is advantageous for its cost-efficiency, since it balances computational resources and accuracy. In particular, it leverages LF data to capture the general trend of the function being regressed, while HF data is used to improve the prediction accuracy. NARGP is built upon traditional co-kriging methods [30,31,45] by recursively combining two *Gaussian Processes* (GPs): the first trained on the LF data, the second on the HF residuals. Each GP assumes a joint Gaussian distribution defined by a mean function and

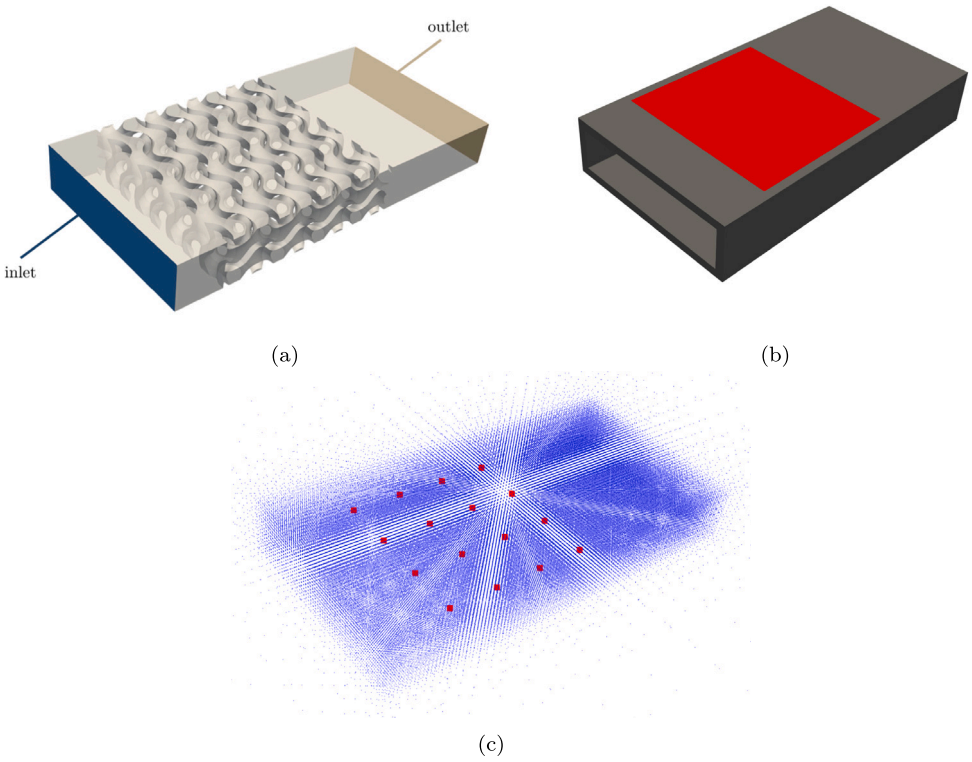

Fig. 7. Fluid domain (a), and aluminium solid case (b) of the HF HX model. In (b), the adiabatic walls are in black, while the heat sources are highlighted in red. Equivalent homogenised HX model (c). In (c), the blue dots are used to indicate the cell centers of computational cells, while the red markers represent the RBF control nodes where the permeability and resistance coefficients are computed by the regression model.

Table 1
Comparison between the HF and homogenised model results for three different mass flow rates. Pressure drop Δp, maximum temperature of the solid phase $T_{(s)}^{max}$, and outlet-averaged temperature of the fluid phase $T_{(f)}^{out}$ are reported in each column with the respective error.

$\dot{m}$ [kg/s]	model	Δp [Pa]	$e_{(\Delta p)}$ [%]	$T_{(s)}^{max}$ [K]	$e_{(T_{(s)}^{max})}$ [%]	$T_{(f)}^{out}$ [K]	$e_{(T_{(f)}^{out})}$ [%]
0.0162	homogenised	96.32	2.87	303.72	0.33	297.52	0.13
0.0162	HF	99.16	–	304.72	–	297.92	–
0.0402	homogenised	463.56	6.18	300.76	0.05	295.02	0.18
0.0402	HF	494.12	–	300.91	–	295.55	–
0.0587	homogenised	895.35	9.52	299.65	0.06	294.97	0.05
0.0587	HF	989.61	–	299.83	–	295.12	–

a covariance kernel (*Matérn* 3/2 *kernel* in this work) with hyperparameters optimised via *Maximum Likelihood Estimation* (MLE) [28] through gradient-based methods. Since GP inference requires the storage in memory and inversion of large covariance matrices, it scales poorly with the dataset size. This makes NARGP best suitable for problems with moderate amounts of training data, where this approach can significantly enhance the accuracy while remaining computationally feasible.

7.1. Application to the equivalent porous media coefficients

The dataset for the *off-line* training of the multi-fidelity model consists of data from two fidelity levels. LF data are represented by the characterisation performed on individual cells provided by periodic BCs as described in Section 5. HF data are represented by the values of K and C_F extracted from the HF simulations described in Section 6. The HF dataset represents a subset of the LF dataset. This is not a strict requirement of the NARGP model, albeit it is advantageous in some cases. To differentiate between the effect of solid walls and head pressure losses, we train a different regression model for each family of lattice cells. These models are queried during the online phase based on the location of the lattice cell in the porous matrix. The regression model for the heat transfer coefficient consists of a simple RBF interpolation and was trained using only LF data. As will be shown later, predictions provided by this simple regression model are accurate enough.

We performed CHT simulations of an entire HX to validate the homogenised model with the new multi-fidelity regressor. The lattice matrix is again composed of a $4 \times 4 \times 1$ array of gyroids with side length $L = 1$ cm and wall thickness 1.5 mm. The HF simulation setup is the one described in Sections 5 and 6, with the embedding of heat transfer. The numerical domain is the one depicted in Fig. 7 (a - b). The fluid phase is governed by the steady incompressible RANS equations with the k–ω SST turbulence model, while the steady heat equation is solved in the solid domain. BCs for the fluid and solid are the same as in Section 5. The $\Delta p°$ applied across the HX are 100, 500, and 1000 Pa. HF simulation results were compared to those of the equivalent homogenised model. For a fair comparison, the resultant mass flow rate from the HF simulations was imposed as an inlet condition in the homogenised model. The resulting pressure drop from the homogenised model was then compared to the pressure drop imposed to the HF runs, following the same methodology used in Section 5. The coefficients

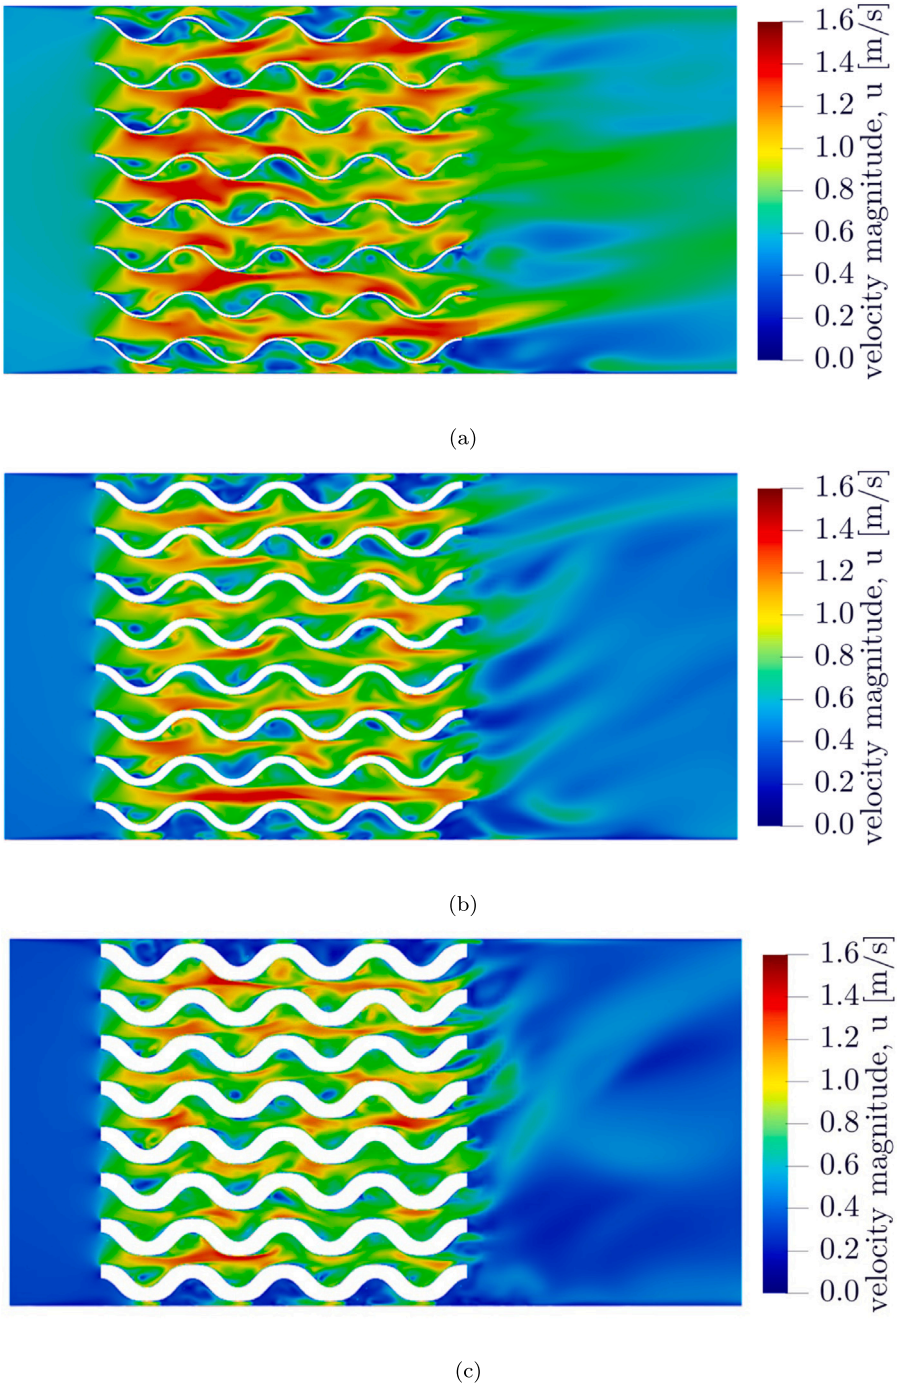

Fig. 8. Velocity magnitude plotted on the mid-plane for an imposed total pressure drop of 2000 Pa and wall thickness of (a) 0.25, (b) 0.75, and (c) 1.5 mm.

K, C_F, and $h_{(sf)}$ used in the homogenised model are predicted by the multi-fidelity model presented in this Section. Due to the non-linear dependence of these coefficients on the local Darcy velocity, fixed-point iterations were used during the momentum solution procedure.

With the improved NARGP multi-fidelity model, the agreement between the models is good. A quantitative comparison is presented in Table 1, which reports Δp, $T_{(f)}^{out}$, and $T_{(s)}^{max}$ for both models. The relative error in pressure drop has a maximum of approximately 10% only when a large pressure drop of 1000 Pa is applied across the HX, as shown in Fig. 11. The maximum solid temperature and the average outlet fluid temperature show excellent agreement, with relative errors below 0.4%.

8. Results and discussion

To validate the multi-scale approach, a test case provided by Rolls-Royce Plc, a potential end-user, was selected. It consists of the optimisation of a HX used in the cooling system for high-power electronics. The geometry is shown in Fig. 12: the grey area represents the solid case, while six hot plates on the top and bottom of the MCHX act as heat sources (shaded red). The inlet and outlet ports are located on the left side of the HX (top and bottom, respectively). The shaded green volume represents the porous matrix. To reduce head pressure losses and facilitate perfusion, each port has a rectangular section, covering about half the HX height. The selected material is aluminium,

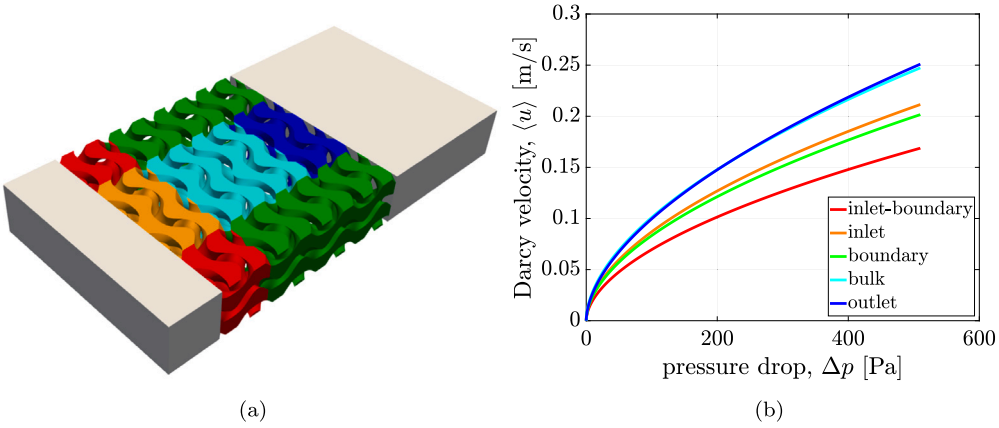

Fig. 9. (a) Distribution of the unit cell types across the HF MCHX simplified model. (b) Exponential relationship between Darcy velocity $\langle u \rangle$ and pressure drop Δp for a gyroid unit cell with a wall thickness of 1.5 mm, located in different regions of the lattice matrix. Solid walls and head pressure are responsible for a reduction in the permeability coefficient K, and larger values of the Forchheimer's coefficient C_F.

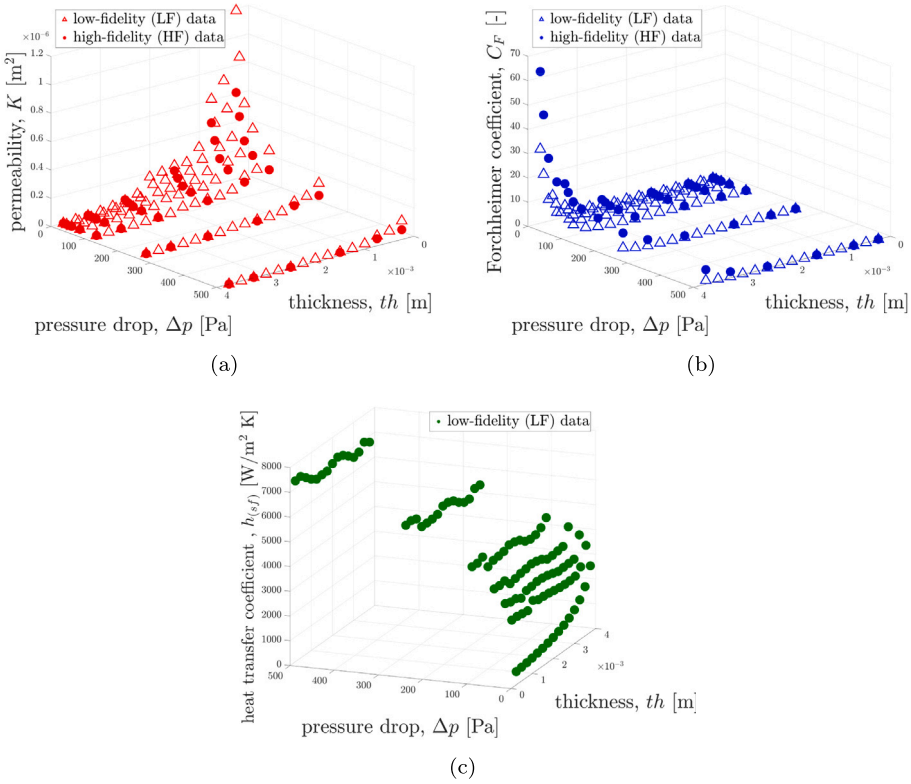

Fig. 10. Comparison between HF and LF data for the permeability K (a) and the Forchheimer coefficient C_F (b), for a gyroid cell located near the inlet and close to a boundary wall. The HF permeability is significantly lower (average relative error equal to 31.06%) than the value obtained from LF simulations. The opposite is true for Forchheimer's coefficient, which HF data are significantly higher than the LF counterpart (average relative error equal to 57.43%). These results suggest that the flow encounters a higher resistance due to head pressure losses and wall boundary effects. The heat transfer coefficient distribution over gyroid wall thickness th and pressure drop Δp is shown in (c).

with thermal properties already reported in Section 5. The 30 RBF control points where thicknesses are optimised are marked in red. A preliminary analysis showed that the optimiser tends to increase wall thickness (i.e., reduce channel width) along the centerline of the porous matrix. This is because narrowing (or widening) channels in that region increases (or decreases) the coolant path length, trading off pressure drop and heat exchange. Hence, a denser placement of RBF nodes is used near the centerline and inlet/outlet.

Given the N_C cell centers of the computational domain porous matrix $\{\mathbf{x}_i\}_{i=1}^{N_C}$, a subset of N_{RBF} cell centers are used as RBF control nodes $\{\mathbf{x}_j\}_{j=1}^{N_{RBF}}$. The wall thicknesses at the control nodes $\boldsymbol{\tau} = [\tau_1, \dots, \tau_{N_{RBF}}]^T \in \mathbb{R}^{N_{RBF}}$ are chosen as design variables. The thickness is then interpolated in each computational cell center using standard RBF interpolation, i.e.

$$\mathbf{th} = \mathbf{W}\,\boldsymbol{\tau} \tag{13}$$

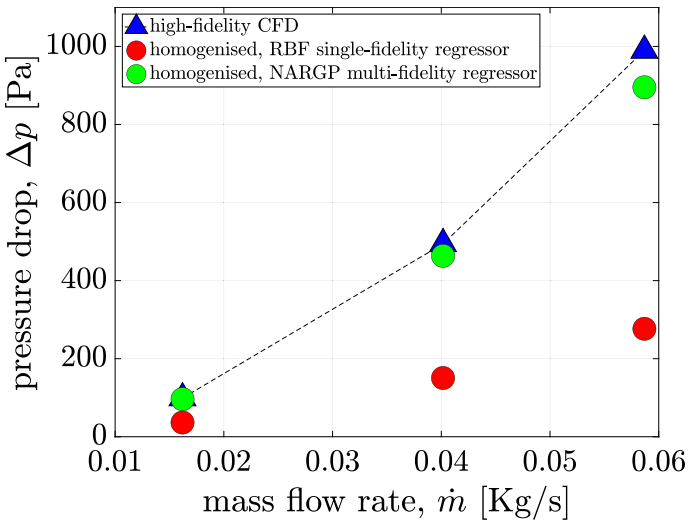

Fig. 11. Inlet-outlet pressure drops Δp, and mass flow rates $\dot{m}$ computed with HF CFD (blue triangles), the homogenised model with RBF single-fidelity regressor (red circles), and the homogenised model with NARGP multi-fidelity regressor (green circles). For each mass flow rate, the pressure drop is underestimated by the homogenised model which only relies on LF data. On the other hand, the homogenised HX with NARGP multi-fidelity model shows excellent agreement in terms of the resulting pressure drop compared to the HF CFD.

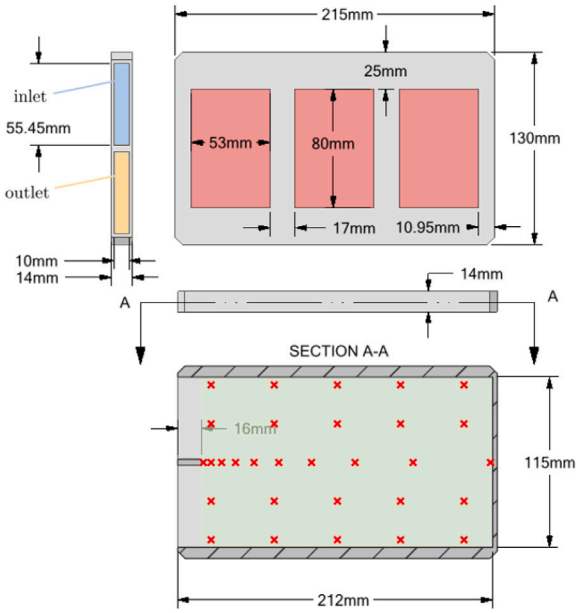

Fig. 12. Schematic depiction of the heat sink test case provided by Rolls-Royce Plc. The red markers on the porous lattice region are the RBF control nodes in which the thicknesses are optimised.

where $\mathbf{W}$ is the RBF weights matrix, and $\mathbf{th} \in \mathbb{R}^{N_C}$ is the vector containing the gyroid wall thickness in each cell center of the porous domain. Once the $\mathbf{th}$ vector is known, a macro-scale CHT run is performed on the MCHX: the equivalent porous media parameters for each cell center $\{\mathbf{x}_i\}_{i=1}^{N_C}$ are estimated by means of the *multi-fidelity* model as a function of the local lattice parameters (gyroid wall thickness distribution $\mathbf{th}$) and flow conditions. The cell-wise equivalent porous media parameters can be written in the form of the vector $\boldsymbol{\psi}(\mathbf{th}) = \{K_i(\mathbf{th}), C_{F_i}(\mathbf{th}), h_i(\mathbf{th})\}_{i=1}^{N_C}$, while the discrete state unknowns of the CHT problems are $\mathbf{y} = [\langle \mathbf{u} \rangle, p, \langle T_{(f)} \rangle, \langle T_{(s)} \rangle]$. By calling $\mathbf{R}(\cdot,\cdot)$ the discretised residual operator of the governing equations and their boundary conditions, the CHT problem can be expressed as:

$$\mathbf{R}(\mathbf{y}; \boldsymbol{\psi}(\mathbf{th})) = \mathbf{0}, \tag{14}$$

meaning that the goal is to find $\mathbf{y}$ such that the CHT equations (described in Section 3, Eq. (7)) are satisfied for the set of porous medium parameters $\boldsymbol{\psi}$ determined by the gyroid wall thickness distribution $\mathbf{th}$. The results $\mathbf{y}$ of each configuration CHT run are then passed to the optimiser, which accounts for them in terms of objective functions and constraints. The objective functions are defined as follows:

$$J_1(\mathbf{th}) = \max_{\mathbf{x}_i \in \Omega_{(s)}} T(\mathbf{x_i}, \mathbf{th}) = T_{(s)}^{max}(\mathbf{th}) \tag{15}$$

$$J_2(\mathbf{th}) = \overline{p}_{\text{in}}(\mathbf{th}) - \overline{p}_{\text{out}}(\mathbf{th}) = \overline{\Delta p}(\mathbf{th}) \tag{16}$$

being $\Omega_{(s)}$ the solid phase domain, and $\overline{p}_{\text{in}}(\mathbf{th})$ and $\overline{p}_{\text{out}}(\mathbf{th})$ the inlet- and outlet-averaged pressure, respectively. The objective $J_1(\mathbf{th})$ is minimised in order to maximise the heat absorption from the solid phase. The objective $J_2(\mathbf{th})$ is minimised, thereby reducing the pressure jump

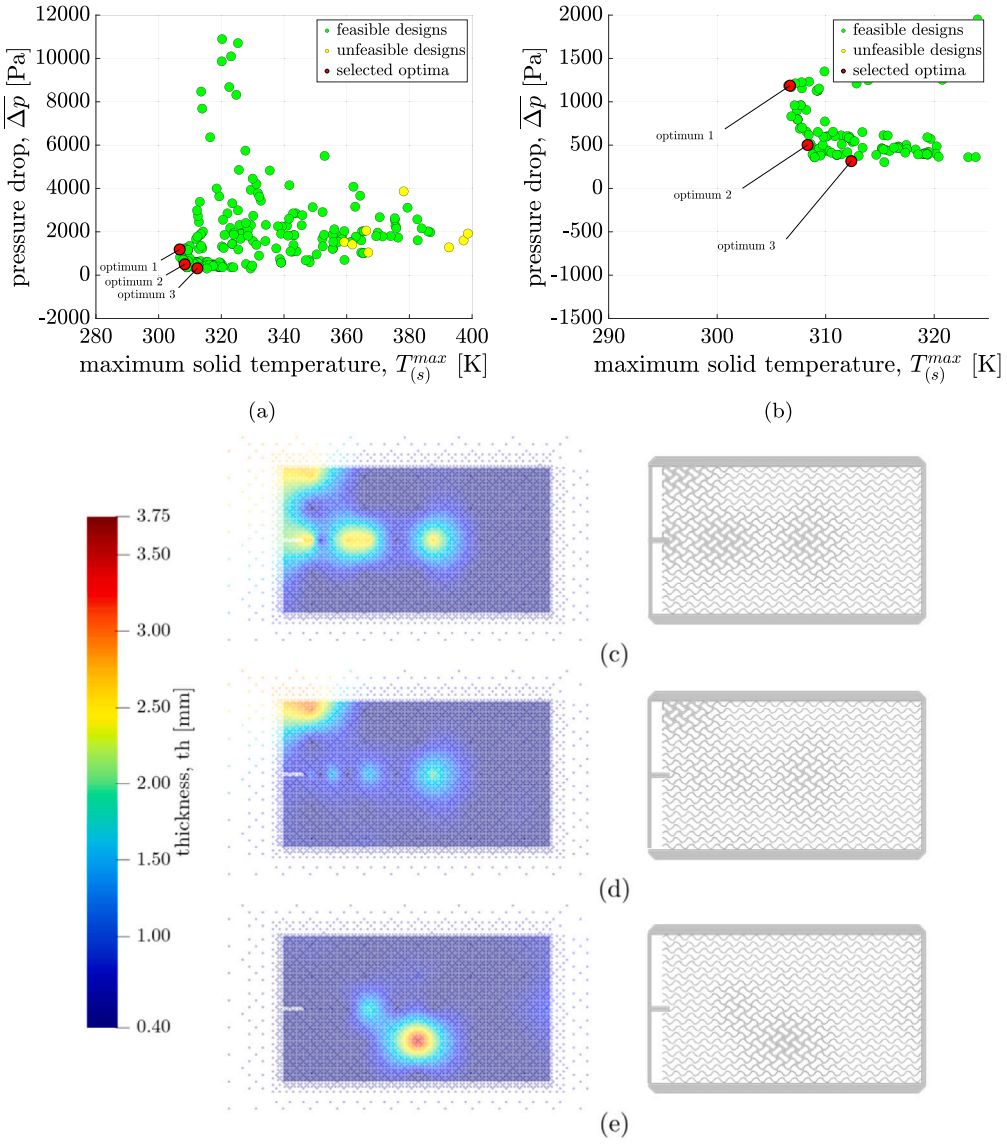

Fig. 13. Overview of the optimisation results and selected designs. (a) Scatter plot of all evaluated configurations during the multi-objective optimisation, showing feasible (green), unfeasible (yellow), and selected optimal designs (red). (b) Zoomed-in view of the Pareto front highlighting the three selected optima. (c–e) Wall thickness distribution and resulting geometry for the three optimised configurations: optimum 1 (c), optimum 2 (d), and optimum 3 (e). For each configuration, the left panel shows the gyroid wall thickness distribution on the mid-plane, while the right panel displays a mid-plane section of the corresponding MCHX geometry.

across the device. Constraints are applied to the maximum solid phase temperature $J_1(\mathbf{th})$, inlet-outlet pressure drop $J_2(\mathbf{th})$, module of the maximum solid phase temperature gradient $\|\nabla T_{(s)}^{max}\|(\mathbf{th})$, and to the local gyroid wall thickness th_i for $i = 1, \dots, N_C$, so that the complete optimisation problem is here formulated:

$$
\begin{aligned}
\min_{\mathbf{th}\in\mathbb{R}^{N_{RBF}}} \quad & \Big(J_1(\mathbf{th}),\ J_2(\mathbf{th})\Big) \\
\text{s.t.} \quad & \mathbf{R}\big(\mathbf{y};\ \boldsymbol{\psi}(\mathbf{th})\big) = \mathbf{0}, \\
& J_1(\mathbf{th}) = T_{(s)}^{\max}(\mathbf{th}) \ \le\ 396\ \text{K}, \\
& J_2(\mathbf{th}) = \overline{\Delta p}\,(\mathbf{th}) \ \le\ 3.5\times10^4\ \text{Pa}, \\
& \|\nabla T_{(s)}^{max}\|(\mathbf{th}) \ \le\ 1,273\ \text{K/m}, \\
& th_{\min} \ \le\ th_i \ \le\ th_{\max}, \quad \forall i.
\end{aligned}
\tag{17}
$$

For $th < 0.4$ mm, distortions occurred during cool-down due to instability, while $th > 3.75$ mm caused poor post-treatment effectiveness. Moreover, the gyroid cell size is fixed at 1 cm to prevent collapse during printing. For these reasons, $th_{min} = 0.4$ mm and $th_{max} = 3.75$ mm are the manufacturer constraints for the local gyroid wall thickness. Objective functions and constraints in Eq. (17) reflect design specifications from Rolls-Royce Plc. All designs assume a uniform initial temperature $T_0 = 294$ K in both fluid and solid domains. The coolant is water, with previously reported thermo-physical properties. The inlet mass flow rate is $\dot{m} = 0.06$ kg/s. Heat input is given through a uniform flux $q_w = 94{,}500$ W/m^2 on the heating plates, while all other walls are adiabatic. The outlet pressure is $\overline{p}_{out} = 0$ Pa.

The optimisation was carried out in modeFRONTIER using the proprietary HYBRID algorithm (NSGA-II coupled with AFilterSQP) in self-initialising mode. A maximum of 500 design evaluations was set. However, the Pareto front was already well populated after approximately 320 evaluations, which was therefore taken as a practical convergence criterion. The initial population consisted of 10 configurations generated through *Latin Hypercube* sampling, where the design variables were the gyroid wall thicknesses at the 30 RBF control nodes.

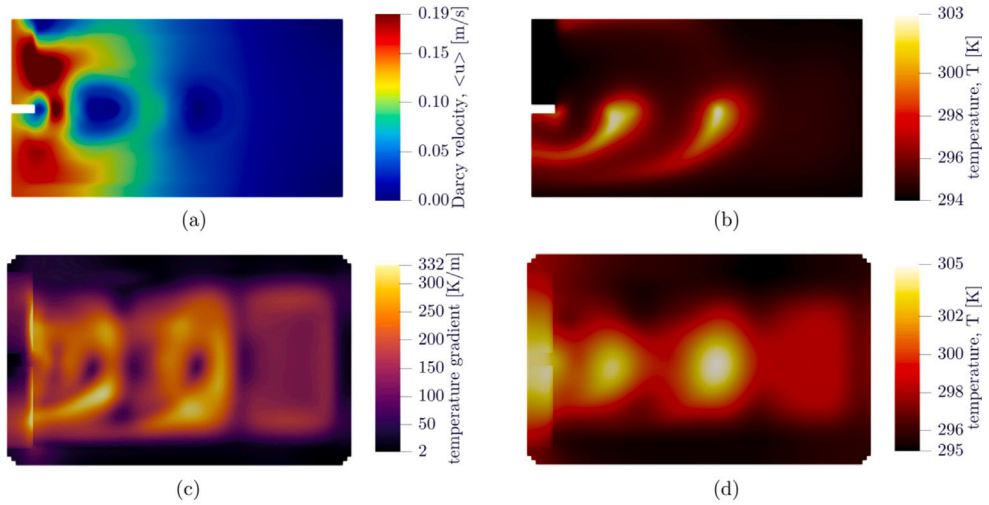

Fig. 14. Numerical results plotted on the mid-plane of the configuration "optimum 1". (a) Magnitude of the Darcy velocity $\langle u \rangle$, (b) temperature in the fluid phase $T_{(f)}$, (c) temperature gradient $\nabla T_{(s)}$ in the solid-phase, (d) temperature in solid-phase $T_{(s)}$.

Table 2
Summary of the adopted optimisation setup parameters with HYBRID algorithm.

HYBRID algorithm optimisation setup parameters	
Number of evaluations	500
Crossover probability	0.90
Mutation probability	1.00
Crossover distribution index	20.00
Mutation distribution index	20.00
Approximate derivative with	RBF approximation
Finite difference relative perturbation	1.00×10^{-7}
Finite difference minimum perturbation policy	Constant
Constant minimum perturbation	1.0×10^{-7}
Range fraction minimum perturbation	0.01
Maximum size of training set	500
Fraction of SQP-generated designs	0.25
Random generator seed	1.00

During the run, four concurrent design evaluations were executed in parallel to improve computational efficiency. The adopted HYBRID setup parameters are reported in Table 2.

Three optimal designs are selected from the Pareto front (Fig. 13 (a - b)): temperature-optimal (optimum 1), pressure-optimal (optimum 3), and an intermediate solution (optimum 2). In the heat transfer optimised configuration, the optimiser increases the wall thickness at the porous matrix centerline to create a high-impedance path that blocks direct inlet-outlet flow. This increases coolant residence time and enhances heat transfer, at the cost of higher pressure losses, as seen in optimum 1 (Fig. 14). Moreover, the high gyroid wall thickness area, which can be noticed in the top left corner of Fig. 13(c), is responsible for the choking of the inlet area, thus increasing the fluid velocity in that region, with the effect of pushing the fluid deeper inside the porous matrix. Again, this produces an increase in the fluid flow path, which has a positive effect on heat transfer due to the increased residence time, but has a negative effect on the pressure drop.

Conversely, in pressure-optimal cases, the optimiser reduces flow impedance. Along the Pareto front, the centerline wall thickness decreases progressively in pressure-optimised designs. In these, tortuosity and recirculation features are sacrificed to favour direct paths and minimise hydraulic resistance, as seen in optimum 2 (Fig. 15) and optimum 3 (Fig. 16). The top-left corner of optimum 2 is again characterised by a large wall thickness (Fig. 13(d)). However, the wall thickness along the centerline becomes less pronounced, and, as a consequence, the impedance of the path which connects directly the inlet to the outlet is decreased, which has a positive effect on the pressure drop (reduced by 58%). At the same time, since the residence time of the coolant is reduced, the dissipated heat power is reduced by 23%. When moving towards the bottom-right part of the Pareto front (where the optimum 3 is located), the optimised configurations exhibit improved performances in terms of pressure drops, but the worst heat transfer. In the configuration corresponding to optimum 3 (Fig. 13(e)), the centerline *septum* is completely absent. Instead, a single region characterised by a large wall thickness in the bottom-central part appears. This thickness distribution allows the fluid to rapidly flow from the inlet to the outlet channel, thus reducing the pressure drop by 75% and 37% with respect to optimum 1 and optimum 2, respectively. For the above-mentioned reasons, also the dissipated heat power $\dot{Q}$ is smaller if compared to previous configurations: 33% and 13% of decrease with respect to optimum 1 and optimum 2, respectively.

Uneven distribution of the flow field inside the porous matrix also results in a temperature maldistribution characterised by strong temperature gradients. This is an undesirable effect as large temperature gradients lead to localised stresses due to the uneven thermal expansion of the material. In the optimised configurations, the maximum temperature gradient is observed in optimum 3. It is located in the solid phase below the heating plates, and it is equal to 636.57 K/m. Thermal gradients are also quite strong at the edges of the six heating plates. Other areas in which thermal gradients are particularly strong can be observed in the transition areas between the inlet/outlet channels and the lattice matrix in all three configurations. The quantitative comparison of results is reported in Table 3.

In order to highlight the improvement brought by the optimisation process, we compare each of the three optimised configurations with a MCHX with a uniform infill lattice composed of gyroids with a wall thickness of 0.4 mm (see Fig. 17). As reported in Table 3, this configuration exhibits a slightly lower pressure drop compared, for instance, to optimal design 3. This is because the coolant directly flows from the inlet to the outlet (lowest impedance path). This, in turn, reduces quite significantly the residence time of the coolant, and ultimately has a negative effect on the exchanged power. The pressure drop for the uniform lattice is approximately 15 Pa lower (4.85%) than the pressure drop observed in optimal design 3. However, due to the shortest residence time of the coolant, the exchanged thermal power is decreased by 28.89% compared to optimal design 3. This result is coherent with the behaviour observed along the Pareto front. Given that pressure drop and exchanged heat power are competing objectives, progression along the front reveals a trade-off. Uniform

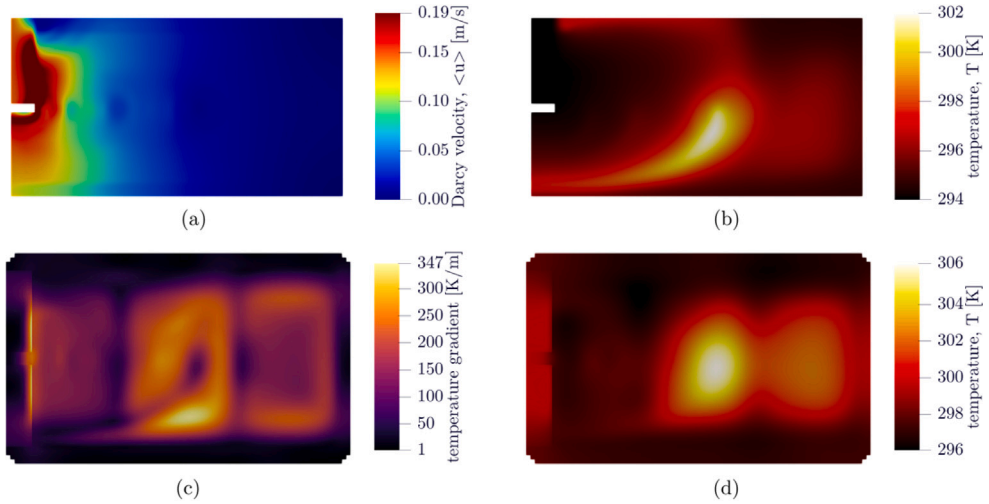

Fig. 15. Numerical results plotted on the mid-plane of the configuration "optimum 2". (a) Magnitude of the Darcy velocity $\langle u \rangle$, (b) temperature in the fluid phase $T_{(f)}$, (c) temperature gradient $\nabla T_{(s)}$ in the solid-phase, (d) temperature in solid-phase $T_{(s)}$.

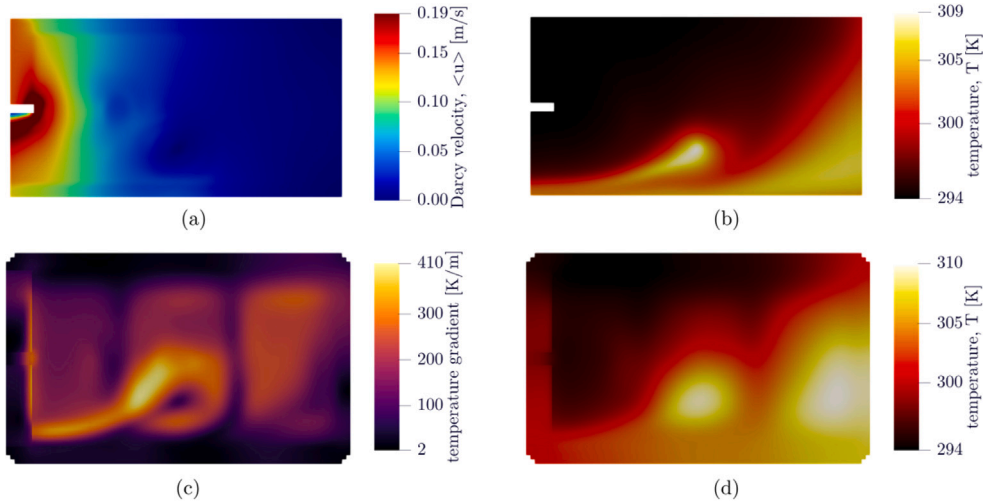

Fig. 16. Numerical results plotted on the mid-plane of the configuration "optimum 3". (a) Magnitude of the Darcy velocity $\langle u \rangle$, (b) temperature in the fluid phase $T_{(f)}$, (c) temperature gradient $\nabla T_{(s)}$ in the solid-phase, (d) temperature in solid-phase $T_{(s)}$.

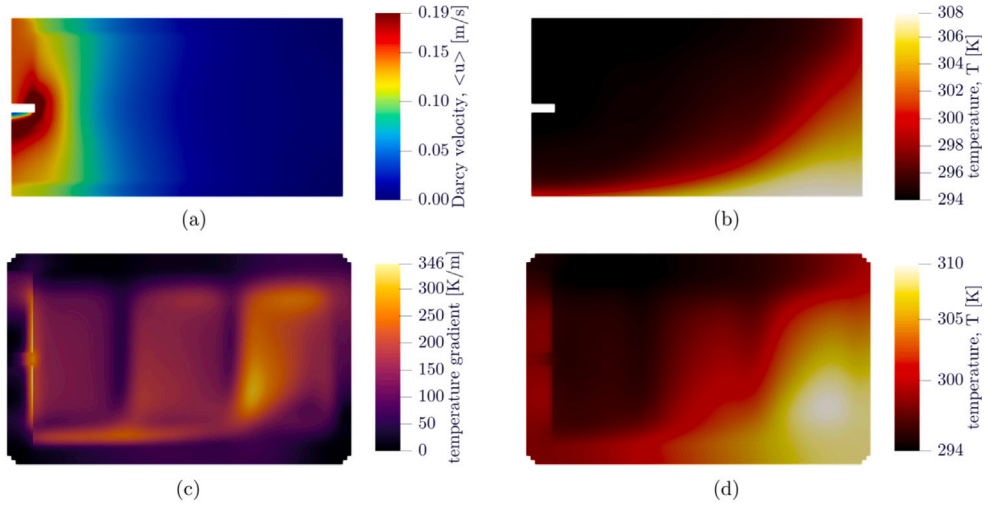

Fig. 17. Numerical results plotted on the mid-plane of the uniform HEX design. (a) Magnitude of the Darcy velocity $\langle u \rangle$, (b) temperature in the fluid phase $T_{(f)}$, (c) temperature gradient $\nabla T_{(s)}$ in the solid-phase, (d) temperature in solid-phase $T_{(s)}$.

Table 3
Comparison of performance indicators for the three optimal configurations obtained from the optimisation, plus the uniform gyroid design.

Performance indicators		optimum 1	optimum 2	optimum 3	uniform
$\dot{Q}$	thermal power exchanged [W]	760.58	588.31	511.63	363.87
$\overline{\Delta p}$	pressure drop [Pa]	1,186.15	501.90	315.30	300.63
$T_{(s)}^{\max}$	max solid temperature [K]	306.71	308.34	312.36	315.83
$\|\nabla T_{(s)}^{max}\|$	max solid thermal gradient [K/m]	614.36	603.05	636.57	573.64
$T_{(f)}^{\text{out}}$	outlet fluid temperature [K]	296.54	295.59	296.33	295.32
$T_{(f)}^{\max}$	max fluid temperature [K]	303.44	301.01	309.32	307.77
$\| \langle \mathbf{u} \rangle^{\max} \|$	max Darcy velocity [m/s]	0.2653	0.3748	0.3493	0.3524

lattice structures tend to yield lower pressure drops but reduced heat exchange, whereas non-uniform lattices result in higher pressure drops yet offer superior heat transfer performance.

9. Conclusions

In this work, we propose a novel methodology for multi-scale simulation and optimisation of MCHXs that is significantly faster than conventional HF CFD, which typically requires meshing complex geometries with hundreds of millions of cells at prohibitive cost. Each cell in the domain is modelled as a gyroid, with thickness controlled by the RBF method. By optimising the RBF weights, the HX geometry is reconstructed. A key contribution of this work is the use of porous media approximation to couple micro- and macro-scales. The micro-scale fluid dynamics is resolved offline and used to train a ML model that captures the relationships between K, C_F, and $h_{(sf)}$ and local lattice geometry and flow conditions. The *off-line* trained ML then is queried in solving the macro-scale CHT problem, predicting local porous media coefficients and enabling the reconstruction of source terms. A further innovation lies in training the ML model on both LF (isolated gyroids) and HF (gyroids from an entire HX) CFD runs, enabling it to account for boundary effects and head losses, thus improving the predictive accuracy. This combination of ML, multi-fidelity data, and offline micro-scale resolution offers a favourable balance between accuracy and cost, enabling optimisation within realistic industrial timelines. Some modelling aspects remain unaddressed in the current workflow, as discussed below.

Wall-roughness. The model does not consider surface roughness, which increases wall friction and pressure drop, and may trigger boundary layer instability, worsening losses but enhancing heat transfer. Consequently, pressure drops may be underestimated. This effect could be included during characterisation by adding a suitable roughness model.

Lattice network limited to gyroid-type. The ML model is trained on gyroid lattices, but other TPMS structures, e.g., Schwarz Primitive, Schwarz Diamond, Lidinoid, Neovius, and IWP, also offer valuable trade-offs in manufacturability, mechanical strength, permeability, and thermal behaviour. The methodology is general and can incorporate these geometries by enriching the CFD database.

The proposed workflow is characterised by a substantial speed-up of CHT runs for the efficient exploration of the design space. Unlike black-box models, the multi-fidelity framework is physically interpretable. HF data integration improves prediction accuracy, mostly where LF models fall short (e.g., boundary effects and head losses), ensuring close agreement with the HF validation runs results. This approach represents a good compromise between efficiency and accuracy, aligning well with time- and resource-constrained industrial design loops. Several future developments can be implemented, such as incorporating roughness, other TPMS types, and optimisation of additional lattice parameters (e.g., frequency and/or local orientation). It can also be extended to other disciplines (e.g., thermo-elastic analysis). Finally, experimental validation on AM prototypes is a key step towards confirming its industrial applicability.

Glossary

AM	Additive Manufacturing
BC	Boundary Condition
CFD	Computational Fluid Dynamics
CHT	Conjugate Heat Transfer
CNC	Computer Numerical Control
FD	Finite Difference
GP	Gaussian Process
HF	High-Fidelity
HPC	High Performance Computing
HX	Heat Exchanger
LF	Low-Fidelity
ML	Machine Learning
MLE	Maximum Likelihood Estimation
MCHX	Micro Channel Heat Exchanger
MSO	Multi-Scale Optimisation
NARGP	Non-Linear Auto-Regressive Gaussian Process
NSGA	Non-Dominated Sorting Genetic Algorithm
R&D	Research and Development
RANS	Reynolds Averaged Navier–Stokes
RBF	Radial Basis Function
REV	Reference Elementary Volume
SDF	Signed Distance Function
SST	Shear Stress Transport
SQP	Sequential Quadratic Programming
TPMS	Triply Periodic Minimal Surface
TRL	Technology Readiness Level

CRediT authorship contribution statement

A. Chiodi: Writing – original draft, Methodology, Investigation, Formal analysis, Data curation. **A. Alaia:** Writing – review & editing, Supervision, Project administration, Conceptualization. **E. Lombardi:** Writing – review & editing, Software. **M. Cisternino:** Writing – review & editing, Software. **K. Gkaragkounis:** Writing – review & editing, Software. **A. Ferrero:** Writing – review & editing, Supervision. **S. Shahpar:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Alessandro Chiodi reports financial support was provided by Ministero dell'università e della ricerca (MIUR). Shahrokh Shahpar reports financial support was provided by UK Research and Innovation (UKRI). Alessandro Alaia, Edoardo Lombardi, and Marco Cisternino report financial support was provided by NextAir EU-Grant Agreement. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Acknowledgements

The authors would like to acknowledge Rolls-Royce Plc for their support and permission to publish this work. This work has been conducted within the framework of the NEXTAIR project jointly funded by the European Union's Horizon Europe research and innovation program, under grant agreement No 101056732, and by the UK Research and Innovation (UKRI), under the UK government's Horizon Europe funding guarantee No. 10038896. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

The work of Alessandro Chiodi is funded by the Italian Ministry of University and Research under the National Recovery and Resilience Plan – PNRR (Ministerial Decree 352/2022), with reference to funding from EU and NextGenerationEU.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ijheatmasstransfer.2025.128035.

Data availability

Data will be made available on request.

References

[1] Kaltra GmbH, Microchannel heat exchangers: Advantages and benefits, 2025, URL https://www.kaltra.com/microchannel-heat-exchangers. (Accessed 18 June 2025).

[2] M. Xu, Y. Wang, J. Li, A review of microchannel heat exchangers and their applications, Am. J. Sci. Technol. 10 (2) (2023) 45–53.

[3] I. Mudawar, M. Hasan, S. Shaaban, Review of Microchannel Heat Sinks for High Heat Flux Removal, Tech. Rep. 2021-03, Purdue University, Cooling Technologies Research Center, 2021.

[4] T. Dixit, Al-Hajri, M.C. Paul, P. Nithiarasu, S. Kumar, High performance, microarchitected, compact heat exchanger enabled by 3D printing, Appl. Therm. Eng. 210 (2022) 118339, http://dx.doi.org/10.1016/j.applthermaleng.2022.118339.

[5] Y. Zhou, S. Shen, T. Liu, P. Li, F. Duan, Effective heat conduction evaluation of lattice structures from selective laser melting printing, Int. J. Heat Mass Transfer 218 (2024) 124790, http://dx.doi.org/10.1016/j.ijheatmasstransfer.2023.124790.

[6] N. Lebaal, A. Settar, S. Roth, S. Gomes, Conjugate heat transfer analysis within in lattice-filled heat exchanger for additive manufacturing, Mech. Adv. Mater. Struct. 29 (10) (2022) 1361–1369, http://dx.doi.org/10.1080/15376494.2020.1819489.

[7] H. Ahmed, H. Sadat, S. Nasrazadani, High-fidelity conjugate heat transfer simulation of micro-channel heat exchanger, J. Adv. Res. Fluid Mech. Therm. Sci. 103 (1) (2023) 81–95.

[8] Y. Luo, Y. Sun, W. Liu, X. Zeng, G. Zhang, Optimization of microchannel heat sinks with flexible vortex generators using hybrid modeling, Eng. Appl. Artif. Intell. 127 (2025) 107583, http://dx.doi.org/10.1016/j.engappai.2024.107583.

[9] V. Glazar, A. Trp, K. Lenic, M. Kirincic, Optimization of air-water microchannel heat exchanger using response surface methodology, Therm. Sci. Eng. Prog. 19 (2020) 100616, http://dx.doi.org/10.1016/j.tsep.2020.100616.

[10] E. Daifalla, S. Shahpar, I. Tristanto, M. Carta, Multidisciplinary optimization of gyroid topologies for a cold plate heat exchanger design, J. Eng. Gas Turbines Power 146 (12) (2024) 121028, http://dx.doi.org/10.1115/1.4055452.

[11] R. Zamolo, L. Bacer, D. Miotti, E. Nobile, M. Munerato, RBF-FD meshless simulation of 3D fully developed flow and heat transfer in triply periodic minimal surfaces, Int. J. Heat Mass Transfer 242 (2025) 126798, http://dx.doi.org/10.1016/j.ijheatmasstransfer.2024.126798.

[12] M. Modrek, K.A. Khan, M.I. Ali, R.K. Abu Al-Rub, Multi-objective topology optimization and numerical investigation of heat sinks based on triply periodic minimal surface lattices, Case Stud. Therm. Eng. 63 (2024) 105255, http://dx.doi.org/10.1016/j.csite.2024.105255.

[13] Y. Jiang, J. Hu, S. Wang, N. Lei, Z. Luo, L. Liu, Meshless optimization of triply periodic minimal surface based two-fluid heat exchanger, Comput.-Aided Des. 162 (2023) 103554, http://dx.doi.org/10.1016/j.cad.2023.103554.

[14] G. Allaire, Shape Optimization by the Homogenization Method, Springer, 2002.

[15] L.V. Gibiansky, S. Torquato, Linking the elastic and conductive properties of composite materials, J. Mech. Phys. Solids 44 (7) (1996) 1115–1142, http://dx.doi.org/10.1016/0022-5096(96)00021-9.

[16] D.M. Kochmann, K. Bertoldi, Exploiting microstructural instabilities in solids and structures: From metamaterials to structural transitions, Appl. Mech. Rev. 69 (5) (2017) 050801, http://dx.doi.org/10.1115/1.4037473.

[17] G. Amhalhel, P. Furmański, Problems of modeling flow and heat transfer in porous media, J. Power Technol. 85 (2011) 55–88.

[18] P.H. Holst, K. Aziz, Transient three-dimensional natural convection in confined porous media, Int. J. Heat Mass Transfer 15 (1) (1972) 73–90.

[19] H. Beji, D. Gobin, Influence of thermal dispersion on natural convection heat transfer in porous media, Numer. Heat Transf. Part A 22 (1992) 487–500.

[20] A. Dybbs, R.V. Edwards, A new look at porous media fluid mechanics – Darcy to turbulent, in: Fundamentals of Transport Phenomena in Porous Media, Martinus Nijhoff, 1984, pp. 199–256.

[21] F.E. Teruel, Rizwan-uddin, Characterization of a porous medium employing numerical tools: Permeability and pressure-drop from Darcy to turbulence, Int. J. Heat Mass Transfer 52 (25) (2009) 5878–5888, http://dx.doi.org/10.1016/j.ijheatmasstransfer.2009.07.017.

[22] D. Kochkov, J.A. Smith, A. Alieva, Q. Wang, M.P. Brenner, S. Hoyer, Machine learning–accelerated computational fluid dynamics, Proc. Natl. Acad. Sci. 118 (21) (2021) http://dx.doi.org/10.1073/pnas.2101784118.

[23] R. Vinuesa, S.L. Brunton, Enhancing computational fluid dynamics with machine learning, Nat. Comput. Sci. 2 (6) (2022) 358–366, http://dx.doi.org/10.1038/s43588-022-00264-7.

[24] Y. Zhao, D.H. Akolekar, J. Weatheritt, V. Michelassi, R.D. Sandberg, RANS turbulence model development using CFD-driven machine learning, J. Comput. Phys. 411 (2020) 109413, http://dx.doi.org/10.1016/j.jcp.2020.109413.

[25] Y. Marioni, P. Adami, F. Montomoli, R. Vazquez-Diaz, S. Sherwin, Machine-learnt turbulence closures for axial compressor cascade with corner separation, in: Proceedings of the European Conference on Turbomachinery Fluid Dynamics and Thermodynamics, 2023, http://dx.doi.org/10.29008/ETC2023-248.

[26] Y. Marioni, P. Adami, R. Vazquez-Diaz, A. Cassinelli, S. Sherwin, Development of machine-learnt turbulence closures for wake mixing predictions in low-pressure turbines, in: Proceedings of ASME Turbo Expo 2022, 2022, http://dx.doi.org/10.1115/GT2022-82531.

[27] J. Hammond, N. Pepper, F. Montomoli, V. Michelassi, Machine learning methods in CFD for turbomachinery: A review, Int. J. Turbomach. Propuls. Power 7 (2) (2022) http://dx.doi.org/10.3390/ijtpp7020016.

[28] C.E. Rasmussen, C.K.I. Williams, Gaussian Processes for Machine Learning, MIT Press, Cambridge, MA, 2006.

[29] P. Perdikaris, M. Raissi, A. Damianou, N.D. Lawrence, G.E. Karniadakis, Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling, in: Proceedings of the Royal Society a: Mathematical, Physical and Engineering Sciences, vol. 473, 2017, 20160751, http://dx.doi.org/10.1098/rspa.2016.0751.

[30] M. Kennedy, A. O'Hagan, Predicting the output from a complex computer code when fast approximations are available, Biometrika 87 (1) (2000) 1–13, http://dx.doi.org/10.1093/biomet/87.1.1.

[31] L.L. Gratiet, J. Garnier, Recursive co-kriging model for design of computer experiments with multiple levels of fidelity, Int. J. Uncertain. Quantif. 4 (4) (2014) 365–386, http://dx.doi.org/10.1615/int.j.uncertaintyquantification.2014006914.

[32] M.P. Bendsoe, O. Sigmund, Topology Optimization: Theory, Methods and Applications, Springer, 2003, http://dx.doi.org/10.1007/978-3-662-05086-6.

[33] ESTECO, modeFRONTIER Optimization Software, 2025, URL https://www.esteco.com/modefrontier. (Accessed 18 June 2025).

[34] A. Turco, MetaHybrid: Combining metamodels and gradient-based techniques in a hybrid multi-objective genetic algorithm, in: Learning and Intelligent Optimisation. LION 2011, in: Lecture Notes in Computer Science, vol. 6683, Springer, 2011.

[35] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II, IEEE Trans. Evol. Comput. 6 (2) (2002) 182–197, http://dx.doi.org/10.1109/4235.996017.

[36] A. Turco, Hybrid – Benchmark Tests, Tech. Rep. 2011-004, ESTECO, 2011.

[37] A. Nakayama, F. Kuwahara, T. Hayashi, Numerical modelling for three-dimensional heat and fluid flow through a bank of cylinders in yaw, J. Fluid Mech. 428 (2004) 139–159.
[38] F. Kuwahara, M. Shirota, A. Nakayama, A numerical study of interfacial convective heat transfer coefficient in two-energy equation model for convection in porous media, Int. J. Heat Mass Transfer 44 (6) (2001) 1153–1159.
[39] Optimad, ImmerFLOW, 2025, URL https://www.optimad.it/. (Accessed 18 June 2025).
[40] Optimad, MIMIC, 2025, URL https://www.optimad.it/. (Accessed 18 June 2025).
[41] OpenFOAM Foundation, OpenFOAM v2012, 2025, URL https://www.openfoam.com/news/main-news/openfoam-v20-12. (Accessed 18 June 2025).
[42] ANSYS Inc., ANSYS Fluent 2022 User's Guide, 2025, URL https://www.ansys.com/resource-library/user-guide/ansys-fluent-user-guide. (Accessed 18 June 2025).
[43] F. Menter, R. Sechner, A. Matyushenko, Best Practice: RANS Turbulence Modeling in ANSYS CFD, Tech. Rep., ANSYS, 2021, pp. 2–95.
[44] Cambridge Flow Solutions Ltd., BOXER mesh: Next-generation meshing technology for CFD, 2025, URL https://www.cambridgeflowsolutions.com/en/products/boxer-mesh/. (Accessed 18 June 2025).
[45] Y. Wu, J. Hu, Q. Zhou, S. Wang, P. Jin, An active learning multi-fidelity metamodeling method based on the bootstrap estimator, Aerosp. Sci. Technol. 106 (2020) 106116, http://dx.doi.org/10.1016/j.ast.2020.106116.