

## Data-driven optimization of rough surfaces for convective heat transfer enhancement

Sanhueza, Rafael Diez; Peeters, Jurriaan W.R.

**DOI**

[10.1016/j.ijheatmasstransfer.2025.127313](https://doi.org/10.1016/j.ijheatmasstransfer.2025.127313)

**Publication date**

2025

**Document Version**

Final published version

**Published in**

International Journal of Heat and Mass Transfer

**Citation (APA)**

Sanhueza, R. D., & Peeters, J. W. R. (2025). Data-driven optimization of rough surfaces for convective heat transfer enhancement. *International Journal of Heat and Mass Transfer*, 251, Article 127313. <https://doi.org/10.1016/j.ijheatmasstransfer.2025.127313>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.



# Data-driven optimization of rough surfaces for convective heat transfer enhancement

Rafael Diez Sanhueza<sup>a,\*</sup>, Jurriaan W.R. Peeters<sup>a</sup>

<sup>a</sup> Process and Energy Department, Delft University of Technology, Leeghwaterstraat 39, 2628 CB Delft, The Netherlands

## ARTICLE INFO

### Keywords:

Heat transfer enhancement  
Dimpled surfaces  
Machine learning  
Direct numerical simulation  
GPU acceleration

## ABSTRACT

Dimpled surface designs are known to be effective at enhancing convective heat transfer. However, optimizing these surfaces can be challenging due to the large parameter space created by the different combinations between geometrical features. In this paper, we combine a machine learning framework with a GPU-accelerated DNS solver to quickly assess the performance of a very large number of surface configurations, and to identify optimal designs. Our neural network can be trained to predict 2-D images with the local Nusselt numbers of rough surfaces within a few hours (in a single GPU), based on their original height maps. During evaluation, our neural network coupled with our parameterized geometrical formulation can evaluate one million dimpled surface designs in less than 45 min using a 64-core CPU architecture; with a low RAM memory footprint per core. Moreover, the GPU-accelerated DNS solver can calculate the Nusselt number of a rough surface within a few hours as well. The study considers a diverse parameter space including dimples with multiple depth profiles, major radiuses, corner effects, and inclination angles. To predict optimal designs, a basic reinforcement loop is created. In the first stage, only randomly chosen dimpled surface designs are selected as training data. The Nusselt numbers for each design are extracted from Direct Numerical Simulations (DNS), performed by the GPU-accelerated turbulent flow solver. Then, a convolutional neural network is trained, and different surface designs in our parameter space are evaluated. In order to advance the reinforcement learning loop, additional DNS cases are run for the optimal predicted surface, and other closely related geometrical variations. After adding these new DNS cases to the training set, the neural network is re-trained, and the process is repeated. Starting from the first iteration of the reinforcement learning loop, our results shows that machine learning can predict remarkably optimized dimpled surface designs, with high Nusselt numbers verified through DNS. Moreover, we find that machine learning chooses dimple configurations that enhance the interaction between roughness elements, even if other dimples with shorter radius (and equal depth) have more heat transfer area. The optimal surface has elongated dimples with opposite inclination angles, which create a zig-zag pattern for the flow near the walls. Additionally, we have shown that at different Reynolds numbers, the optimal geometry is different as well. We analyze other plausible optimal dimpled surface designs within our parameter space, and we find that machine learning correctly identified the adequate parameters to maximize heat transfer. Therefore, we conclude that machine learning is a highly effective tool to identify optimized designs for convective heat transfer enhancement.

## 1. Introduction

Turbulent flows past rough surfaces can be found in different engineering applications. In most cases, rough surfaces tend to increase the drag resistance of transportation systems, leading to higher energy losses and fuel consumption. However, certain categories of rough surfaces can produce favorable results, such as enhancing heat transfer while only producing a modest increase in pressure losses [1,2]. Therefore, it is possible to design special patterns for rough surfaces that maximize heat transfer inside engineering equipment. Maximizing

the mean Nusselt number  $\overline{Nu}$  of a rough surface is important for engineering applications, since heat transfer equipment can be made smaller overall. Different families of rough surfaces have been studied to accomplish this goal, such as riblets [3] or dimples [4]. Under special conditions, it has been shown that it is even possible to increase the net ratio between the Stanton number ( $\overline{St}$ ) and the skin friction factor ( $\overline{C_f}$ ) of a rough surface [2,4]. An extensive review of rough surface patterns for heat transfer enhancement can be found in [4–6]. In many studies, it has been noted that dimples have a high potential to increase heat

\* Corresponding author.

E-mail address: [R.G.DiezSanhueza-1@tudelft.nl](mailto:R.G.DiezSanhueza-1@tudelft.nl) (R.D. Sanhueza).

## Nomenclature

### Operators

—	Average over entire surface
+	Dimensionless value (scaled by $Re_\tau$ )

### Sub-indexes

0	Smooth wall reference (same $Re_b$ )
---	--------------------------------------

### Variables

$\alpha$	Dimple global inclination
$\theta$	Angular coordinate inside the dimples
$C_f$	Skin friction factor
$C_{a-b}$	Corner effects for dimples
$D$	Depth profiles for dimples
$d_{ref}$	Reference depth for dimples
$DNS$	Direct numerical simulation
$H$	Height function
$L$	Length of the computational domain
$N$	Number of grid points
$Nu$	Nusselt number
$Pr$	Prandtl number
$Q$	Heating load
$R$	Dimple radius
$St$	Stanton number
$W_p$	Pumping power
$x$	Streamwise direction
$y$	Wall-normal direction
$z$	Spanwise direction

transfer in a system [6–10]. Moreover, dimpled surfaces are easy to manufacture, and different modifications can be considered, such as elliptical elements, rotated shapes [11], vortex generators [12,13]. Other studies have further considered elliptical protrusions [14], trapezoidal geometries [15], staggered arrangements [16], teardrop shapes [17], tubes with large dimples or protrusions [18,19].

One challenge while optimizing dimpled surface designs is that the parameter space created by the combinations between all geometrical features is very large. It is well-known that modifications in the dimple shapes can lead to substantially different results [4,6]. Therefore, employing machine learning is necessary to assess the heat transfer performance of different dimple shapes. The current study uses a previously developed neural network architecture [20], which is trained using existing DNS data, and it is capable of scanning the height map of a rough surface, and then predicting the local distribution of Nusselt numbers  $Nu(x, z)$  or skin friction factors  $C_f(x, z)$ . While traditional convolutional neural networks must process an entire image for each predicted scalar quantity, our neural network architecture [20] can recycle the results of intermediate convolutional layers, and thus it is capable of generating 2-D maps with all  $Nu(x, z)$  predictions after processing an image only once. Theoretically, this reduces the time and space complexity of the problem from quadratic to linear complexity, during both the training and evaluation stages. In our practical application, we indeed observe that the wall clock times for the training and evaluation of the neural network are reduced by several orders of magnitude. An example showcasing our convolutional neural network architecture, together with the filtering methodology previously developed [20], can be found in this Github repository [21].

Thanks to the efficiency of the machine learning system, it becomes possible to create a reinforcement-learning study, where millions of

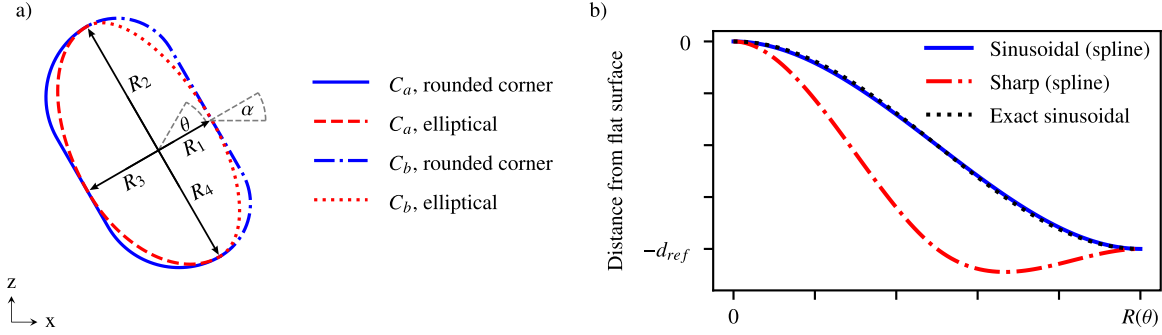
rough surface designs are evaluated in a short time searching for optimal designs. Within this context, one important advantage of machine learning over gradient-based algorithms like the discrete adjoint method, is that the impact of large discrete changes in surface features can be assessed directly, without relying on infinitesimal gradients. Estimating the gradients with respect to changes in surface features is computationally expensive [22], especially for DNS with large grid sizes. Moreover, gradient-based optimizers also risk running into other numerical issues, such as convergence to local minima close to the starting configuration. About the usage of other optimization methods, we highlight that in Section 3, our current machine learning framework is able to find optimized dimpled surfaces using only 20 training samples, which is less than the number of degrees of freedom in the system. Classical optimization methods would require hundreds, or even thousands, of surface design evaluations (through DNS) before converging.

To further strengthen our computational framework, we use a GPU-accelerated DNS solver written in Fortran, allowing for fast and accurate generation of high-fidelity data. The GPU-based DNS solver is used both for generating training data, and for verifying the heat transfer performance of selected dimpled surface designs. In general, GPU-based DNS solvers can be substantially faster than even multicore CPU algorithms [23]. Therefore, this work combines for the first time the neural network architecture developed by Sanhueza et al. [20] with a GPU-accelerated DNS solver to generate optimized dimple surface designs. This paper is organized as follows: in Section 2, the methodology of the study will be presented, along with the parameter-space created for the rough surfaces. Then, in Section 3, the results of the study are described, which is followed by the conclusions in Section 4.

## 2. Methodology

### 2.1. Geometrical variations

During the current study, a diverse collection of dimpled surface designs is considered, including variations in different types of relevant features. A schematic representation of the geometrical variations available for every dimple can be found in Fig. 1. In this scheme, different shape modifications are highlighted, such as the radial profiles ( $R_{1-4}$ ), dimple curvature ( $C_{a-b}$ ), inclination angle ( $\alpha$ ), and depth profile ( $D$ ). The values considered for each of these parameters are listed in Table 1. Regarding the depth profiles shown in Fig. 1(c), each spline is internally controlled by an anchor point located at  $R(\theta)/2$ . The first profile (blue) is a nearly sinusoidal shape, which resembles a classical dimple study, whereas the second profile (red) has a sharper inclination angle. Either of these shapes can be chosen during the machine learning optimization study. As a side-note, more shape variations were initially considered. However, during a preliminary study, intermediate depth profiles only produced gradual variations in our subsequent DNS results, and the optimum Nusselt number  $\bar{Nu}$  was found for one of the two extreme depth profiles (round/sharp). Regarding the dimple curvature effects shown in Fig. 1(b), the “elliptical” shape is obtained by considering four different ellipses, with major radiuses  $R_{1-4}$  spaced at  $90^\circ$  intervals. In contrast, the “circular corner” effect uses the smaller dimple radius to draw a circular arc at  $90^\circ$ . From a physical perspective, this circular corner effect can increase the windward area of the dimples, where the maximum heat transfer typically occurs. However, this profile could also have the opposite effect, and expand the region with recirculation inside the dimples, which is detrimental for heat transfer. Therefore, the variations in the curvature profiles pose an interesting challenge for the machine learning study. The inclination angle ( $\alpha$ ) is physically relevant as well, since inclined dimples can enhance vortex generation compared to spherical dimples [24,25]. However, the machine learning system must determine the optimal inclination angle ( $\alpha$ ), and how different dimple shapes will affect vortex interactions. Finally, regarding the depth of the dimples  $d_{ref}$ , this parameter is kept at




**Fig. 1.** Schematic representation of the parameters considered for the dimples. The sketch (a) highlights the location of the radiuses  $R_{1-4}$  within the dimples and their inclination angle ( $\alpha$ ) with respect to the  $x$ - $z$  plane. The legend of subplot (a) indicates the different curvature effects available for the dimples. Finally, the subplot (b) shows the two possible depth profiles for the dimples. Within this plot,  $d_{ref}$  is the reference depth for the dimples, whereas  $R(\theta)$  is the local dimple radius at any given angle  $\theta$ . Please note that  $R(\theta)$  is interpolated from subplot (a) taking into consideration the local curvature effects. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

**Table 1**

Parameters considered for the dimpled surfaces. The variable  $\alpha$  is the inclination angle for the dimples, whereas  $R_{1-4}$  are the four major radiuses for each dimple. The parameters  $C_{a-b}$  are the two different curvature variations, and  $D$  is the depth profile for every dimple. The influence of the previous parameters in the dimple geometry is sketched in Fig. 1.

Variable	Alternatives						
$\alpha$	-45°	-30°	-15°	0°	15°	30°	45°
$R_{1-4}$	30%	50%	70%				
$C_{a-b}$	Elliptical	Rounded corner					
$D$	Rounded	Sharp					



**Fig. 2.** Representation of the staggered dimple arrangement for the simulations. The shape of each dimple (A/B) is given by the alternatives listed in Table 1.

a fixed value of  $d_{ref} = \delta/10$ , where  $\delta$  is the half channel height:  $\delta = L_y/2$ . If  $d_{ref}$  is changed, our DNS analysis showed that the Nusselt number  $\overline{Nu}$  tends to increase monotonically, even for very large  $d_{ref}/\delta$  ratios. Such DNS cases require more grid cells, and the identification of adequate initial conditions to avoid divergence. Good initial conditions are often obtained by running preliminary DNS cases with less pronounced geometrical features. However, the process of running additional simulations is computationally expensive, and it requires the formulation of empirical selection criteria. Therefore, all DNS simulations were performed only for  $d_{ref} = \delta/10$ . As a general note, our dimpled surfaces with  $d_{ref} = \delta/10$  can display Nusselt numbers up to 53% higher than smooth wall channel flows under equal pressure losses (at  $Re_\tau = 180$ ), and the simulations can start from almost any initial condition.

Regarding the global dimple arrangements, the final study consists in optimizing two rows of staggered dimples that are periodically repeated, as is shown in Fig. 2. Each row of dimples has a unique shape, according to the alternatives listed in Table 1. The total number of configurations available for the system is approximately 3.5 million, after removing repeated entries due to periodicity and physical symmetry in the spanwise direction. The dimples within each row (in the  $z$ -direction) have identical shape, since a preliminary study showed that the flow across the dimples mainly interacts with other elements aligned in the streamwise direction. Moreover, adding more dimple variations in the spanwise direction (per row) exponentially increases

the number of combinations available, and it was not found the change the outcome of preliminary ML optimization studies. Regarding the spacing between the dimples, all shapes are packed as closely as possible, while preserving the staggered grid arrangement. The center of every dimple was chosen to be the average position of the region with a position deeper than 75% of the nominal depth ( $d_{ref}$ ). This formula does not change the center of uniform dimples, and it ensures that the deepest parts of asymmetrical dimples follow a staggered arrangement. This is important to reduce flow obstruction, since the deepest area of the dimples has a large impact in flow circulation, and the regions with the highest amount of heat transfer are usually located in front of them. Therefore, the current formula ensures that the dimples have a good alignment within our dataset, although other methods could be tested in the future.

## 2.2. Machine learning framework

The overall framework of the machine learning study corresponds to a basic reinforcement-learning loop, as it is shown in Fig. 3. In this framework, a DNS database is generated first (3.A), containing purely random dimpled surfaces. Then, a neural network is trained (3.B) to predict the local Nusselt numbers  $Nu(x, z)$  distributed across the rough surfaces. Using the newly trained neural network, the mean Nusselt numbers  $\overline{Nu}$  are predicted for all possible dimpled surface combinations (3.C) according to Table 1. Considering the best performing dimpled surface (3.D), new, yet closely related designs are generated randomly, according to the sub-steps of (3.E). Additional channel flow cases are run for each of these new dimpled surfaces, and their results are added to the DNS database (3.F), in order to repeat the reinforcement-learning loop. Beyond the framework presented in Fig. 3, it is important to note that the size of the neural network trained in step (3.B) was increased starting from the third iteration of the reinforcement learning loop. This modification was necessary to fit the increasing amount of training data available, since the size of our DNS database approximately triples by the third stage of the reinforcement learning study. Otherwise, the neural network would display a relatively high bias even when performing predictions for the training dataset. The details of the neural network architecture are discussed at the end of this sub-section.

In steps (i-iii) of Fig. 3.E, the methodology to choose variations of the optimal dimpled surface configuration predicted by machine learning is described. In this methodology, the dimple configurations are changed, by picking one parameter, and giving it random values for both dimples. The main benefit of this approach is that the influence of each physical feature is tested separately. If more parameters were changed at the same time, many configurations would have a distorted shape, which is physically far from the optimal predicted

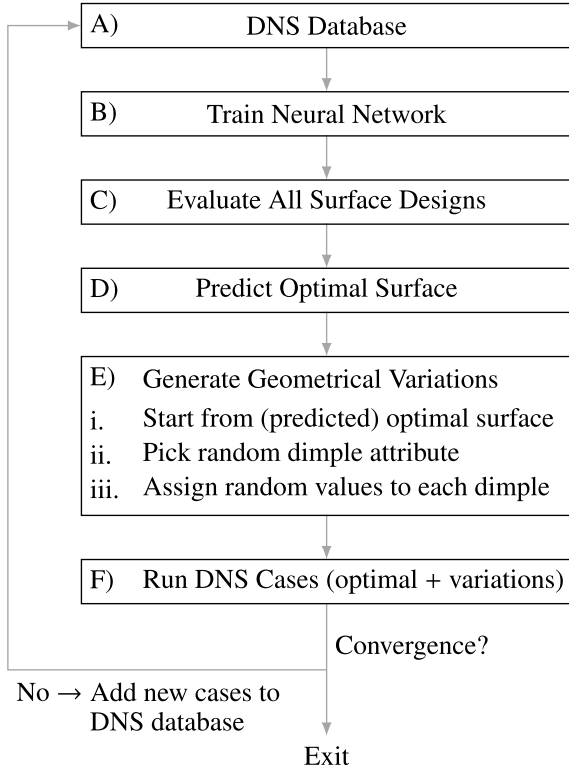


Fig. 3. Steps of the basic reinforcement-learning loop for the current study.

configuration. For example, the dimples configurations can be highly sensitive to changes in their depth profile, or their radii. Therefore, the current methodology establishes a better trade-off between exploring new design characteristics, and keeping a shape that is reasonably related to the best predicted configuration. The total number of possible variations is 93, according to Table 1. Only 20 of these random dimple configurations are sampled in every stage of the reinforcement learning loop. Therefore, the machine learning system must still use sparse information to make predictions at any given stage.

Regarding the details of the neural network architecture, the system is based on the optimized convolutional neural network (CNN) architecture described in Sanhueza et al. [20]. This machine learning system is able to scan the height map of a rough surface  $H(x, z)$  and predict its local Nusselt number distribution  $Nu(x, z)$  with linear time complexity, during both the training and evaluation stages. In our target application, this is thousands of times faster than a traditional convolutional neural network, which must process an entire image for every scalar quantity predicted. Predicting the local  $Nu(x, z)$  values can be beneficial, because it forces machine learning to build a predictive model that accounts for local flow effects. Moreover, this methodology also provides a one-to-one ratio between the input data ( $H(x, z)$ ) and the predicted quantities ( $Nu(x, z)$ ) for the neural network, which otherwise would be very unbalanced if  $\overline{Nu}$  was predicted directly. Regarding traditional surface metrics to predict  $\overline{Nu}$ , such as the skewness or the effective slope, these quantities only correspond to averaged values, and they do not take into account the alignment between the roughness elements [26]. During the current study, it was observed that identical roughness elements with optimized alignment could substantially increase heat transfer. Therefore, traditional surface metrics do not form a good basis for the optimization study. Our machine learning framework to predict local  $Nu(x, z)$  distributions is well suited to take into account how local flow effects influence heat transfer.

The overall neural network architecture for the current study is described in Table 2. Here, two neural networks are presented. The

Table 2

Neural network architectures to predict the local Nusselt numbers  $Nu(x, z)$ . The total number of trainable parameters is 4461 for the initial neural network, and 8161 for the larger ML model.

Layer	Input channels	Output channels	Kernel size	Dilation	Activation function
Initial neural network					
DSC	1	20	3	$1 \times 1$	Yes
DSC	20	20	3	$2 \times 2$	Yes
DSC	20	20	3	$4 \times 4$	Yes
DSC	20	20	3	$8 \times 8$	Yes
DSC	20	20	3	$16 \times 16$	Yes
DSC	20	20	3	$32 \times 32$	Yes
Conv. 1D	20	1	1	–	–
Large neural network					
DSC	1	20	3	$1 \times 1$	Yes
DSC	20	20	3	$1 \times 1$	Yes
DSC	20	20	3	$2 \times 2$	Yes
DSC	20	20	3	$2 \times 2$	Yes
DSC	20	20	3	$4 \times 4$	Yes
DSC	20	20	3	$4 \times 4$	Yes
DSC	20	20	3	$8 \times 8$	Yes
DSC	20	20	3	$8 \times 8$	Yes
DSC	20	20	3	$16 \times 16$	Yes
DSC	20	20	3	$16 \times 16$	Yes
DSC	20	20	3	$32 \times 32$	Yes
Conv. 1D	20	1	1	–	–

first network has a reduced number of layers, which is used during the first two optimization stages of the study, due to the smaller amount of training data available in the DNS database. In the third reinforcement learning iteration, a larger neural network is employed, in order to fit the increasing amount of training data. The details of each DSC module mentioned in Table 2 are sketched in Fig. 4. This architecture follows the principles described in [20], although the number of activation functions inside the network was reduced. This minor modification was made to reduce the number of trainable parameters, and because small empirical differences were found in the context of the current study. The neural networks listed in Table 2 have a total of 4461 and 8161 trainable parameters, for the smaller and the larger neural network respectively.

### 2.3. GPU-accelerated direct numerical simulations (DNS)

In this study, the dimpled surfaces are simulated as rough walls within rectangular channel flows. The numerical framework is primarily based on the dimensionless form of the incompressible Navier-Stokes equations:

$$\nabla \cdot \mathbf{u} = 0, \quad (1)$$

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \frac{1}{Re_\tau} \nabla^2 \mathbf{u} + S_f, \quad (2)$$

$$\partial_t T + (\mathbf{u} \cdot \nabla) T = \frac{1}{Re_\tau Pr} \nabla^2 T + S_q, \quad (3)$$

In Eqs. (1), (2), (3), the variables  $\mathbf{u}$ ,  $P$ ,  $T$  refer to the fluid velocity vector, pressure and temperature respectively. The parameter  $Re_\tau$  is the friction Reynolds number of the fluid, whereas  $Pr$  is the molecular Prandtl number. The constants  $S_f$  and  $S_q$  are source terms that induce fluid motion and heat transfer in their respective equations [27–31]. The geometry considered for the study is a planar channel flow with dimpled walls, as it is shown in Fig. 5. In this scheme,  $H(x, z)$  is the height function for the dimpled surface, the  $y$ -coordinate is the wall-normal direction, and  $L_y$  corresponds to the full-channel height. Periodic boundary conditions are considered in the streamwise ( $x$ ) and spanwise directions ( $z$ ). The challenge of the heat transfer optimization






Fig. 4. Schematic representation of the depthwise separable convolution (DSC) modules mentioned in Table 2. The abbreviation *BN* refers to the 2-D batch normalization operations applied before each PReLU activation function.

problem is to maximize the mean Nusselt number ( $\overline{Nu}$ ), which is given by:

$$\overline{Nu} = \frac{\overline{q_w}}{\lambda (T_b - T_w)/L_y}, \quad (4)$$

where  $\lambda = 1/(Re_\tau Pr)$  is the dimensionless conductivity of the fluid,  $T_w$  is the wall temperature, and  $T_b$  is the bulk temperature of the fluid. As it can be observed, maximizing the Nusselt number of the system ( $\overline{Nu}$ ) is equivalent to decreasing the temperature difference ( $T_b - T_w$ ), since all other parameters are fixed in the current study. The mean heat flux  $\overline{q_w}$  is also fixed, because the energy balance must be in equilibrium with respect to the volumetric source term  $S_q$ .

From the global simulation settings, it is important to highlight that the channel flows operate with a constant mean pressure loss  $\Delta P$ , since the friction drag losses at the rough walls are compensated by the momentum source term  $S_f$ . As a consequence, in this study, the Nusselt number for the channel flow is maximized without changing the average pressure losses. This marks a contrast with respect to other formulations, where the highest Nusselt number might be found for configurations with very high pressure losses. Another implication of the current settings is that the skin friction factor  $\overline{C_f}$  is not a direct target of optimization for the study.

In practical terms, our optimization study resembles cases like the design of a heat exchanger. We seek to provide a constant heating load  $Q$  with fixed pressure losses  $\Delta P$ , and the objective is to maximize the Nusselt number  $\overline{Nu}$  to ensure that the system can operate with the lowest temperature difference:  $(T_b - T_w)$ . Here, the drag losses should be measured in terms of the hydraulic pumping power, which is proportional to  $\Delta P U_b$ , among other fixed parameters ( $L_x$ ,  $L_z$ , etc.). Based on these settings, dimpled surfaces with higher hydraulic resistance will have a lower bulk velocity ( $U_b$ ) and Reynolds number ( $Re_b$ ), and hence their pumping power will not be greater. Despite this fact, since the Nusselt number  $\overline{Nu}$  is highly correlated with  $Re_b$ , surfaces with high hydraulic resistance will be subject to a trade-off, where their  $\overline{Nu}$  values are penalized by their lower bulk velocity. Therefore, the results of the optimization study are physically-relevant.

In the literature, most studies tend to address a different challenge, which is maximizing the thermal efficiency  $(Nu/Nu_0)/(C_f/C_{f,0})$  with respect to a smooth wall (sub-index 0); for a fixed bulk Reynolds number ( $Re_b$ ) [4,6]. While surfaces with high thermal efficiency are valuable, this metric is not directly applicable to this study, because the changes in the skin friction factor ( $C_f$ ) will modify the pressure losses  $\Delta P$ , and hence the channel flow will reach another equilibrium with different  $Re_b$ . As a result, we highlight that maximizing  $\overline{Nu}$  is the right optimization target for the current study (with fixed  $\Delta P$  and  $Q$  values).

Regarding the discretization and simulation parameters, the study is performed considering air with  $Pr = 0.71$ ,  $S_f = 1$ ,  $S_q = 1$  and  $L_y = 2$ . The friction Reynolds number ( $Re_\tau$ ) can be either 180 or 395, depending on the simulation settings. The grid size is  $(N_x \times N_y \times N_z) = (800 \times 256 \times 400)$  in each spatial direction for the simulations at  $Re_\tau = 180$ , whereas the  $Re_\tau = 395$  use a grid size of  $(N_x \times N_y \times N_z) = (1920 \times 592 \times 960)$ . The size of the channel flows in the streamwise ( $L_x$ ) and spanwise directions ( $L_z$ ) varies between  $L_x = [5.66, 7.50]$  and  $L_z = [2.82, 4.13]$ , depending on the size of the dimples simulated. These domain sizes are consistent with other works about DNS for turbulent flows past rough surfaces [27,32,33].




Fig. 5. Scheme of the channel flow geometry, with the height function  $H(x, z)$  representing the dimpled surfaces. The wall-normal direction is given by the  $y$ -coordinate, where  $x$  is the streamwise direction. The  $z$ -coordinate, perpendicular to the image, is the spanwise axis. The plane  $y = 0$  is coincident with the average height for the dimpled surface ( $H(x, z)$ ). The variable  $L_y$  is the full channel height, and thus  $y = L_y/2$  is the symmetry plane at the channel center.

In order to simulate the behavior of the rough surfaces, the immersed boundary method (IBM) is used. Our IBM implementation is inspired by the Fadlun scheme [34], although we use a ghost point inside the solid instead of forcing the first (fluid) point above the rough surfaces. Extensive validation benchmarks were performed, obtaining identical results as Peeters and Sandham [27]. Due to the refined grid, all the current DNS cases have dimensionless cell sizes in the range  $\Delta x^+ < 1.69$ ,  $\Delta z^+ < 1.86$ , and a uniform mesh size of  $\Delta y^+ \approx 0.9$  is kept near the rough surfaces. The mesh size is small in the  $x$  and  $z$  directions, since the dimples have a significant curvature, and this helps to reduce the distance at which interpolation points are considered. The grid size further complies with the basic discretization requirements of similar simulations [27–29], which can be expressed as  $(\Delta x^+, \Delta y^+, \Delta z^+) < (5, 1, 5)$ . Other discretization constraints related to the length-scale of the rough surfaces do not apply to the current study, since the dimples are substantially larger than the smallest grid scales. More details about the validation of the GPU-based DNS solver can be found in Appendix.

From a fluid mechanics perspective, it is important to mention that all DNS cases with dimpled surfaces operate in turbulent regimes. Both smooth wall and rough surface cases display turbulent flow behavior at  $Re_\tau = 180$  for moderate roughness heights [27–29].

About the software implementation, our GPU-accelerated DNS solver is written in Fortran, using GPU-aware MPI and OpenACC for cross-platform compatibility. The numerical algorithm uses the fractional-step method, following the methodology used by Peeters and Sandham [27]. Finite difference expressions are used for the spatial discretization of the Navier–Stokes equations, whereas the temporal discretization is given by an Adams–Bashforth scheme. The immersed boundary method uses ghost points along fixed Cartesian directions, as previously mentioned. To fulfill the continuity equation, a pressure–Poisson equation is implicitly solved using spectral FFT-based methods, achieving  $\mathcal{O}(n \log n)$  time complexity for large-scale applications. In the Poisson solver, Fourier transformations are applied in the streamwise ( $x$ ) and spanwise ( $z$ ) directions, whereas the wall-normal direction ( $y$ ) retains its finite difference discretization and a tridiagonal solver is used instead. To ensure high-performance in multi-GPU configurations, the computational domain is divided into 1D slabs along the  $y$ -direction for each MPI task. This implies that all FFT operations along the  $x$ - $z$  directions are local to each GPU device. To solve tridiagonal equations along the  $y$ -direction, we incorporate a parallel solver, based on the parallel cyclic reduction (PCR) algorithm [35,36].

## 2.4. Scalability analysis

Regarding the scalability of our optimization framework, only the running times of the DNS solver are affected by the higher Reynolds number ( $Re_\tau$ ). The input images passed to the machine learning system, with the local Nusselt number distributions  $Nu(x, z)$ , can be rescaled to a lower resolution similar to the DNS grid for  $Re_\tau = 180$ . In our experience, this resolution is enough to capture both the geometrical features of the dimpled surfaces and the changes in the local  $Nu(x, z)$  distributions, even for flows at higher Reynolds numbers. Therefore, the running times of our machine learning framework are constant for flows at different Reynolds numbers. In the case of the DNS solver, the Poisson solver tends to be the main performance bottleneck for large-scale simulations, since the remaining subroutines of the DNS solver either perform halo exchanges or local kernel computations. Despite these challenges, our DNS code incorporating a parallel tridiagonal solver is able to achieve strong scalability (by grid size) when comparing simulations at either  $Re_\tau = 180$  or  $Re_\tau = 395$ . Detailed benchmarks about the scalability of DNS solvers with PCR algorithms can be found in [35–38].

## 3. Results

### 3.1. Optimization for dimples surfaces at $Re_\tau = 180$

In order to initialize the reinforcement-learning loop, described in Fig. 3, a DNS database with completely random dimple variations is generated first. The dimpled surfaces found in this initial database are plotted in Fig. 6. Here, it can be verified that all the 20 DNS cases have unique patterns and different surface topologies. Using the GPU-accelerated DNS solver, the local Nusselt number distributions  $Nu(x, z)$  for each of these 20 random surfaces are calculated. Based on these results, the smallest neural network described in Table 2 is trained. From the 20 DNS cases available, 18 DNS cases are used as training data, and 2 DNS cases are left as a small cross-validation set to monitor overfitting. After training the neural network, approximately 3.5 million different surface designs are evaluated, seeking to maximize the mean Nusselt number  $\overline{Nu}$ . The best three performing configurations predicted by the neural network in the first reinforcement-learning generation are shown in Fig. 7. Here, it can be seen that machine learning predicts that the optimal surface consists of two elongated dimples with opposite inclination angles ( $\alpha$ ). This result is very interesting, since none of the shapes in the training data (Fig. 6) had a similar pattern. The Nusselt numbers predicted by machine learning, and the DNS verification for the top-3 configurations are also shown in Fig. 7. From these results, it can be noted that the ranking assigned to the dimpled surface designs is correct with respect to their Nusselt numbers  $\overline{Nu}$  extracted from the DNS data, even though the neural network slightly under-predicts the mean Nusselt number  $\overline{Nu}$  compared to the DNS results.

For the second iteration of the reinforcement learning loop, new dimpled surfaces are generated, according to the methodology described in steps (i-iii) of Fig. 3. The surfaces created can be found in Fig. 8. Here, it can be noted that the new surfaces contain large variations from a physical perspective, despite only changing one parameter at a time. The neural network is thus re-trained using an extended dataset with 41 DNS cases: the 20 original surfaces shown in Fig. 6, the top-1 ranked surface identified by machine learning (first row in Fig. 7), and the variations shown in Fig. 8. After re-training the neural network, the results shown in Fig. 9 are obtained. The main difference with respect to the first generation of the reinforcement learning loop is that inclination angle of the dimpled surfaces was changed from  $45^\circ$  to  $30^\circ$ . The variations in the top predicted configurations are minor, and they follow the same patterns observed back in Fig. 7.


In the third generation of the reinforcement-learning study, 20 additional surfaces are considered. Now the DNS database contains a total of 62 DNS cases, since the optimal design from the second

generation (Fig. 9) is also added to the database. Due to the increased amount of training data, the size of the neural network was increased. Otherwise, the bias of the neural network can be higher, even for training samples. The results of the machine learning study are found in Fig. 10. Here, it can be seen that the top predicted configurations have a very similar topology to the second reinforcement learning stage, although small variations in the curvature of the perimeter for the dimples are observed. All dimpled surfaces still have high potential for heat transfer, according to the DNS results. However, it is observed that the minor topological changes observed among the top dimpled surfaces do not further influence the outcome of the study, and the reinforcement-learning iterations are concluded.

To investigate which surface features contribute the most to heat transfer enhancement, Fig. 11 presents a comparison between multiple variations from the optimal surface predicted in the second generation of the reinforcement learning study. Here, the geometry of the dimpled surfaces is shown along with their respective Nusselt numbers  $\overline{Nu}$  extracted from DNS data. From this comparison, it can be noted that the configuration predicted by machine learning is substantially better than similar designs. This result indicates that the neural network correctly found an optimized dimpled surface design starting from low-quality random training data, which had little relation with the optimal patterns. Many reinforcement learning frameworks would require hundreds, if not thousands, of iterations before they can converge to an adequate design. In contrast, our machine learning system quickly found correct geometrical parameters, such as the inclination angle  $\alpha = 30^\circ$ , corner profiles, depth and aspect ratio for the dimples among similar alternatives. Regarding these surfaces, it can also be highlighted that the dimples with a circular perimeter, or lower aspect ratio, are more closely packed and have more heat transfer area. Therefore, machine learning took into account the interactions between roughness elements while assessing the potential for heat transfer of different surface designs.

Finally, it is interesting to investigate if the performance of the optimal dimpled surface can also be explained by investigating the local flow structures. Flow patterns are generally important for dimpled surfaces, since most heat transfer takes place in the windward faces, whereas backflow areas (with recirculation) have a negligible contribution. Due to this reason, the averaged Nusselt numbers  $\overline{Nu}$  of dimpled surfaces are greatly dependent on any heat transfer enhancement effects at the windward faces, due to factors such as the shape or alignment of dimple elements. Based on this context, the higher Nusselt numbers for dimpled surfaces with opposite inclination angles can be attributed to the creation of a zig-zag pattern in the flow near the surfaces. This phenomenon is visualized in Fig. 12, where the streamlines for the top surface from the second generation of the reinforcement learning loop (Fig. 9) are compared with the dimples with circular perimeter shown in Fig. 11 and other closely related geometrical variations. Here, it can be clearly observed that the dimples with opposite inclination angles have minimal fluid recirculation, since the streamlines inside each dimple follow a spiral pattern [39], where the flow from previous dimples is deviated resulting in a large heat transfer rate in the frontal faces. In contrast, the circular dimples suffer from significant flow recirculation in their interior [40–43], which is detrimental for local heat transfer. Another reason for the lower thermal performance of circular dimples is that the flow perturbations created by one element do not have a significant impact on the upcoming dimples. This phenomenon lowers the thermal performance of aligned roughness elements in general. Beyond the alignment of the roughness elements, the larger  $\overline{Nu}$  values observed for dimples with sharper corners, or steeper depth profiles, is related to the increased size of the frontal area where the highest local Nusselt numbers  $Nu(x, z)$  are found.

Regarding the influence of the dimple inclination angle, Fig. 12 shows that dimples with a moderate inclination of  $15^\circ$  have a much higher degree of recirculation than the optimal configuration ( $30^\circ$ ).



**Fig. 6.** Dimpled surfaces found in the DNS database for the first generation of the reinforcement-learning framework. The sub-images are drawn at scale to represent their size in physical coordinates:  $(L_x \times L_z)$ .

This is evidenced by the additional loops of the spiral vortices inside these dimples. Dimples with a higher inclination angle of  $45^\circ$  display minimal recirculation, yet it can be observed that the spiral vortices inside the dimples are highly distorted. Finally, the dimples in aligned arrangement at  $30^\circ$  present a drastically different flow pattern, where the flow tends to hover above the dimples, which is less than optimal for heat transfer. In summary, it can be observed that both the inclination angle of the dimples and the staggered arrangement have a large impact in the flow patterns for every configuration, and the resulting heat transfer.

Therefore, our machine learning framework can be a useful tool to predict surface designs that maximize heat transfer, and we provide a framework to iterate using additional results obtained from DNS simulations. Moreover, the local Nusselt numbers  $Nu(x, z)$  predicted for each configuration serve as an additional verification tool, to check whether the distributions follow physically plausible patterns, or if the mean Nusselt number  $\overline{Nu}$  are affected by spurious non-physical oscillations. Further design insights can also be obtained from  $Nu(x, z)$  to predict

optimal heat transfer patterns, or to understand the sensitivity of the neural network to changes in the dimple geometry.

To provide a broader context, the performance metrics reported by Ligrani et al. [4] and the Reynolds analogy factor  $(2St/C_f)$  [44] were computed for the top dimpled surface from the second generation of the machine learning study, and other canonical flow cases. These values are reported in Table 3, along with the dimensionless ratio between heat transfer ( $Q$ ) and pumping power ( $W_p$ ) for each configuration. Here, it can be observed that the top dimple configuration also has a significantly higher ratio  $Q/W_p/(Q_0/W_{p,0})$ . Furthermore, the metric  $\overline{Nu}/Nu_0/(C_f/C_{f,0})^{1/3}$  is also higher for the optimized surface. However, according to the Reynolds analogy factor and the thermo-hydraulic efficiency, the optimized surface performs less well. Thus, we conclude that the optimization method yields an improved ratio of thermal load to pumping power. The latter is the direct result of the Nusselt number being much higher for the top dimpled surface than it is for the circular dimple surface. Moreover, our channel flow simulations have fixed pressure losses  $\Delta P$ , which implies that the bulk Reynolds number  $Re_b$  is subject to change. Therefore, the value reported by



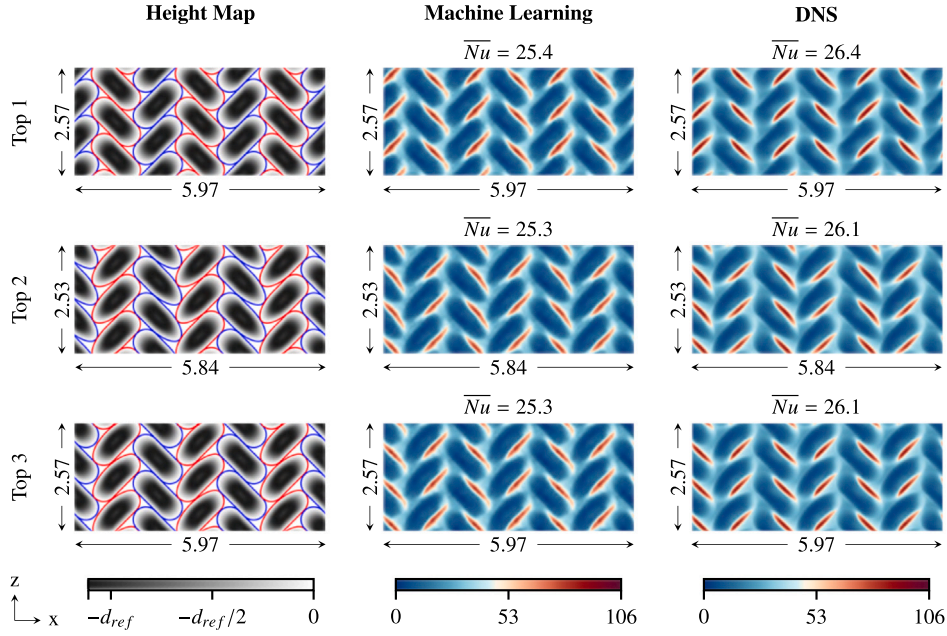


Fig. 7. Top ranked dimpled surfaces found during the first generation of the reinforcement learning loop at  $Re_\tau = 180$ . The bulk flow moves from left to right.

Table 3

Thermal performance metrics for the highest performing dimpled surface found in the second generation of the reinforcement learning study at  $Re_\tau = 180$ , and other closely related alternatives. The sub-index “0” refers to a flat-plate channel case at identical bulk Reynolds number ( $Re_b$ ) as the target configuration.

Case	$\frac{Q/W_p}{Q_0/W_{p,0}}$	$\overline{Nu}$	$C_f$	$\frac{2St}{C_f}$	$\frac{\overline{Nu}/\overline{Nu}_0}{(C_f/C_{f,0})^{1/3}}$	$\frac{\overline{Nu}/\overline{Nu}_0}{C_f/C_{f,0}}$
Top dimpled surface (gen. 2)	1.63	28.0	0.024	1.00	1.54	0.81
Circular dimples (same $Re_b$ )	1.25	17.8	0.014	1.10	1.17	0.88
Smooth walls, $Re_\tau \approx 110$ (same $Re_b$ )	1	13.2	0.0091	1.25	–	–
Smooth walls, $Re_\tau = 180$ (same $\Delta P$ )	–	18.1	0.0079	1.13	–	–

the thermo-hydraulic efficiency does not necessarily reflect the final performance of the system.

Finally, to emphasize that the computed flows at  $Re_\tau = 180$  operate in the turbulent regime, Fig. 13 presents a snapshot of the instantaneous velocity field  $U^+$  at the plane  $z = L_z/2$  for the top-performing dimpled surface from the second generation of the reinforcement learning study. Here, it can be seen that the velocity field displays the similar characteristics as turbulent flow snapshots reported in other studies [27].

### 3.2. Optimization for dimples surfaces at $Re_\tau = 395$

In the previous section, it was shown that machine learning can accurately predict optimized dimpled surface designs for flows at a friction Reynolds number of  $Re_\tau = 180$ . However, to further demonstrate the ability of the optimization method, we repeated the procedure for flows at  $Re_\tau = 395$ . To begin this new phase of the study, a new DNS database was generated, considering the randomized dimpled surfaces previously shown in Fig. 6. Then, the neural network was re-trained using the new DNS database with flows at  $Re_\tau = 395$ , and the evaluation procedure with millions of combinations was repeated. The top configurations predicted by machine learning are shown in Fig. 14. Here, it can be seen that the top-ranked dimpled surface resembles the previous optimal design for  $Re_\tau = 180$ . However, instead of sharper (rounded) corners, elliptical corners are chosen instead. To further investigate this outcome, an additional DNS was performed at  $Re_\tau = 395$ , but with sharper corners instead. The last row of Fig. 14 shows a comparison with respect to a dimpled surface with sharper (rounded) corners. From the DNS results, it can be seen that the elliptical corners are indeed beneficial to increase the Nusselt number ( $Nu$ ) at  $Re_\tau = 395$ .

The streamline patterns for DNS cases with both corner effects can be found in Fig. 15. Here, it can be observed that the elliptical corners create spiral vortices more closely aligned with the dimple geometry, and that most streamlines are ejected together towards the end of each dimple. In contrast, the dimples with sharp (rounded) corners have a more dispersed flow pattern, with the streamlines spilling over the dimples and being ejected at different locations. This implies that the geometry is not closely aligned with the vortices generated, and thus it is natural to expect a higher Nusselt number from the dimples with elliptical corners at  $Re_\tau = 395$ .

A comparison between the performance metrics for the newly optimized surface at  $Re_\tau = 395$  and other relevant cases can be found in Table 4. Here, the Nusselt numbers  $\overline{Nu}$  tend to be higher than before (Table 3), since this quantity grows together  $Re_b$  or  $Re_\tau$  in our physical formulation. However, we highlight that the newly optimized surface (with elliptical corners) has higher values for  $Q/W_p/(Q_0/W_{p,0})$  than the previous study at  $Re_\tau = 180$ . This indicates that the dimpled surface is operating efficiently, achieving both a high Nusselt number and a good ratio of heat transfer to pumping power. Again, the ratio  $\overline{Nu}/\overline{Nu}_0/(C_f/C_{f,0})^{1/3}$  has a favorable value for the optimized surface.

### 3.3. Discussion of results

Regarding the quantitative results presented here, it should be mentioned that the topologies found will only be optimal for the investigated Reynolds numbers. Moreover, the optimized topology was found using Dirichlet boundary conditions. This modeling assumption may not always be correct when dealing with heat transfer equipment. A good example would be a heat exchanger without any multiphase




Fig. 8. Additional dimpled surfaces added to the DNS database for the second generation of the reinforcement-learning framework.

Table 4

Thermal performance metrics for the highest performing dimpled surface during the optimization study with  $Re_\tau = 395$ , and other closely related alternatives. The sub-index “0” refers to a flat-plate channel case at identical bulk Reynolds number ( $Re_b$ ) as the target configuration.

Case	$\frac{Q/W_p}{Q_0/W_{p,0}}$	$\overline{Nu}$	$C_f$	$\frac{2St}{C_f}$	$\frac{\overline{Nu}/\overline{Nu}_0}{(C_f/C_{f,0})^{1/3}}$	$\frac{\overline{Nu}/\overline{Nu}_0}{C_f/C_{f,0}}$
Top dimpled surface ( $Re_\tau = 395$ )	1.78	52.8	0.023	0.88	1.65	0.77
Circular dimples (same $Re_b$ )	1.65	46.8	0.0199	0.9	1.54	0.79
Smooth walls, $Re_\tau \approx 222$ (same $Re_b$ )	1	21.7	0.0073	1.14	1	1
Smooth walls, $Re_\tau = 395$ (same $\Delta P$ )	–	36.6	0.0063	1.16	–	–

phenomena on either the hot or cold side. However, the study’s methodology can be extended to the point where the surface topology is not only optimized on one side, but on both the hot- and cold-side in tandem. In such a case, only the configuration of the DNS solver would need to change to allow for conjugate heat transfer, but the

convolutional neural network would still work and thus, the methods in this paper could be used to optimize topologies that are more relevant to industry. In subsequent studies, the dimple depth should be included in the parametrization of the dimple geometry. Furthermore, the same general methodology can be used for different designs such



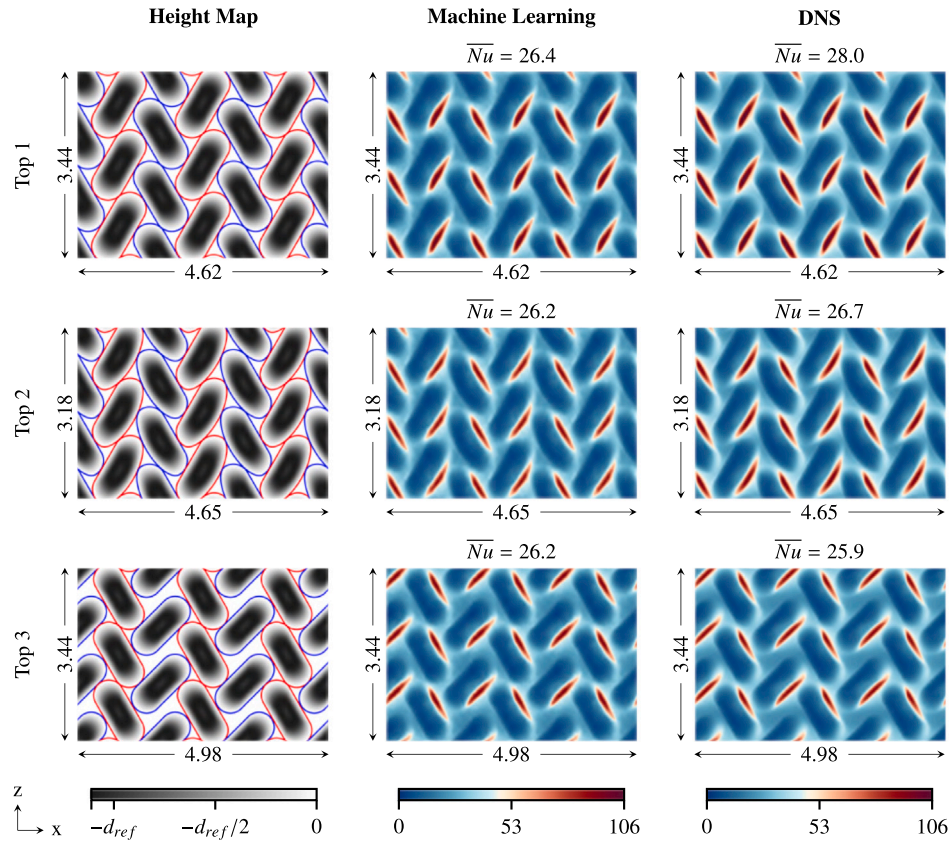


Fig. 9. Top ranked dimpled surfaces found during the second generation of the reinforcement learning loop at  $Re_\tau = 180$ . The bulk flow moves from left to right.

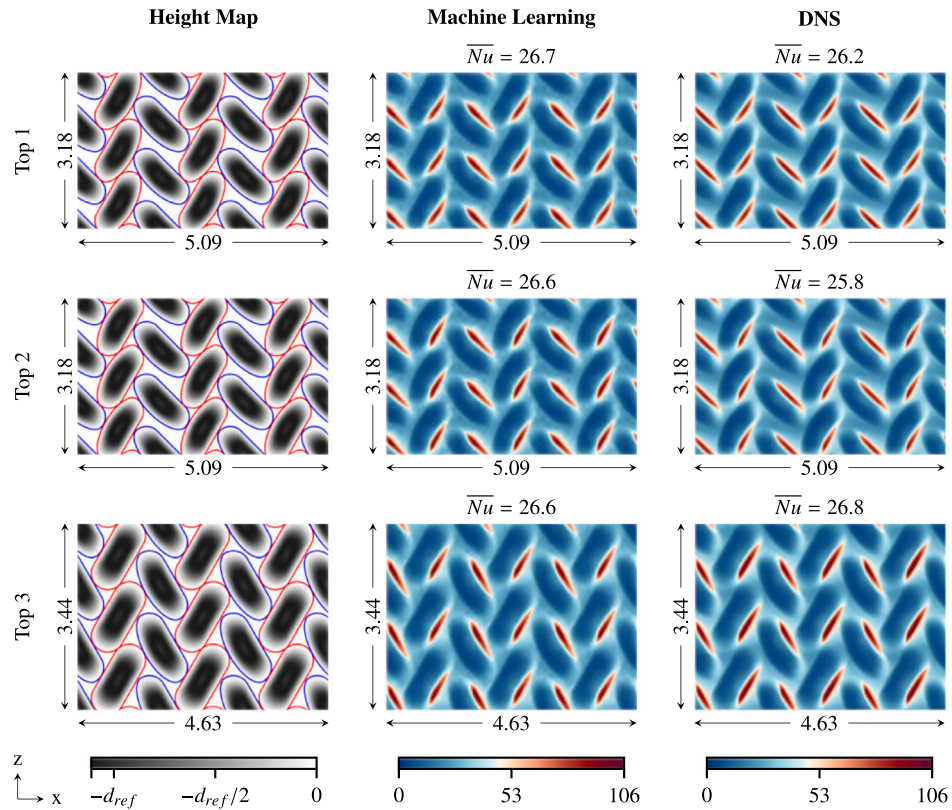
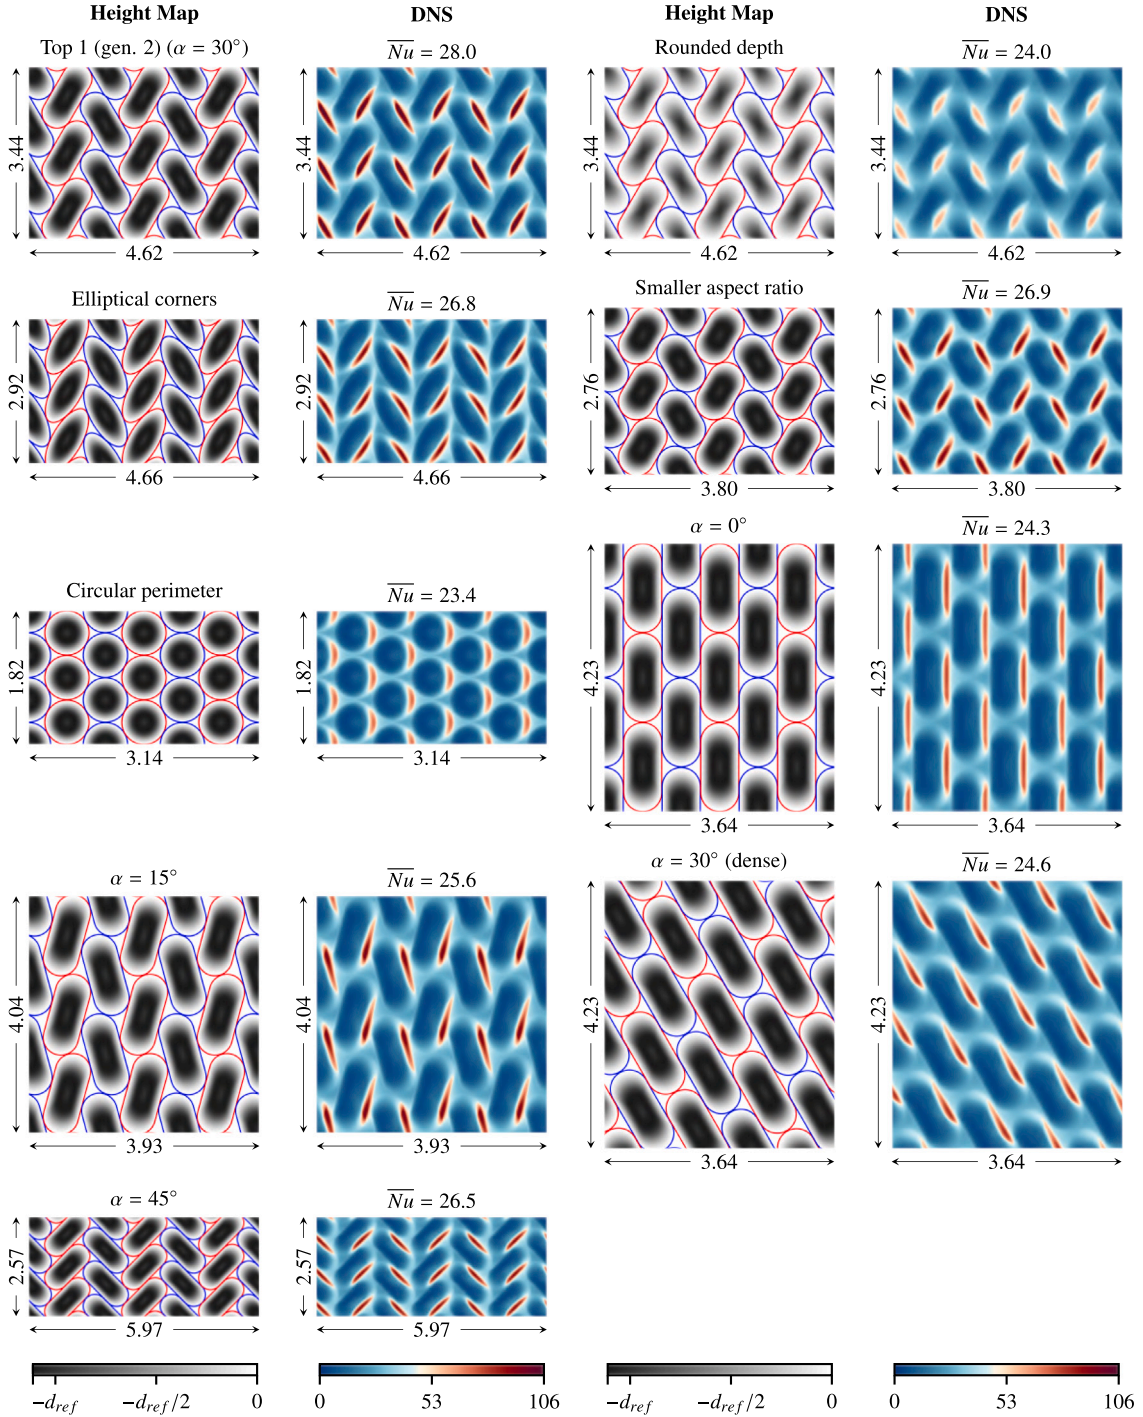


Fig. 10. Top ranked dimpled surfaces found during the third generation of the reinforcement learning loop at  $Re_\tau = 180$ . The bulk flow moves from left to right.



**Fig. 11.** Comparison between the optimal dimpled surface found in the second generation of the reinforcement learning loop with other closely related geometries at  $Re_\tau = 180$ . The bulk flow moves from left to right.

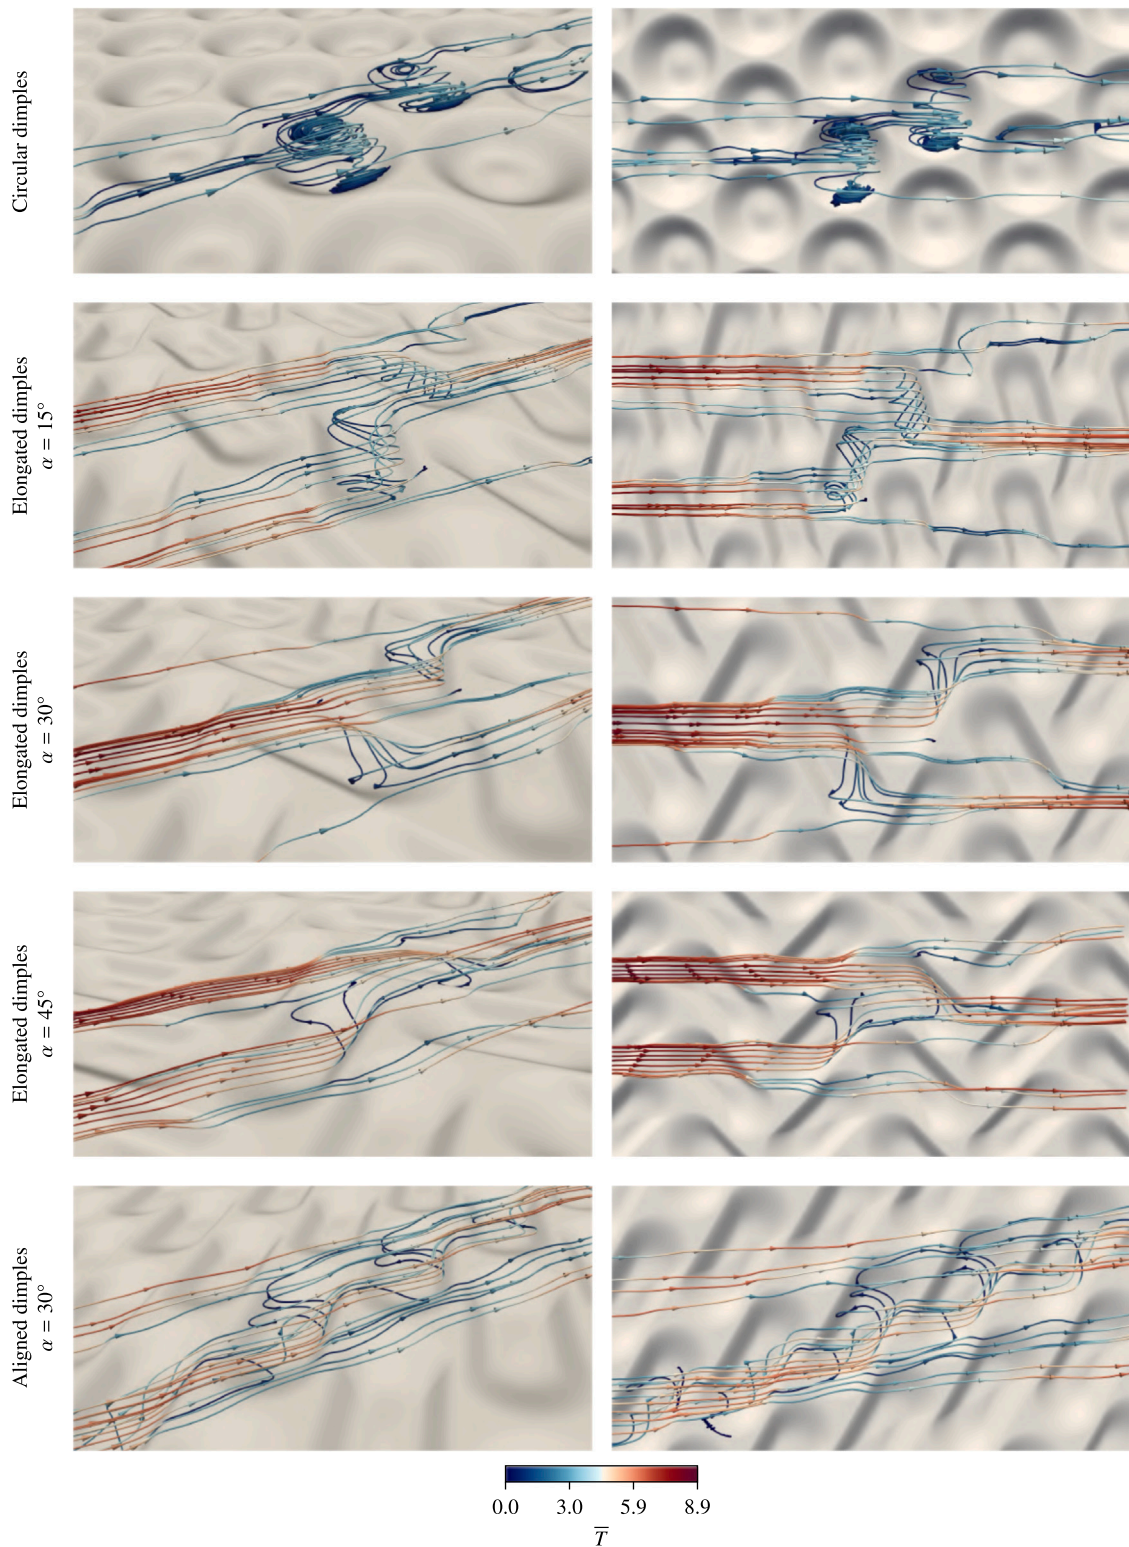
as transverse bars or riblets, as long as the parametrization is changed accordingly. Finally, to include much higher Reynolds numbers, Large Eddy Simulations could be used instead of Direct Numerical Simulations to extend the applicability of the methods to a wide range of industrial cases.

#### 4. Conclusions

In this study, we present a machine learning framework to optimize rough surfaces for convective heat transfer enhancement. The procedure starts by considering a DNS database with purely random

surface designs. Then, a neural network is trained using the existing DNS data, and a new optimal design is predicted within our parameter space. To advance in the reinforcement learning loop, the DNS database is augmented by simulating both the new optimal surface, as well as closely related random variations of this design. The neural network can predict highly optimized dimpled surface designs, starting from the first iteration of the reinforcement learning loop. The rough surfaces identified by machine learning contain elongated dimples with opposite inclination angles, which create a zig-zag pattern for the flow near the walls. This design is highly effective for heat transfer enhancement, and further analysis shows that it is substantially more effective than other






**Fig. 12.** Streamlines for the dimples with circular perimeter (left) shown in Fig. 11, and the top-ranked dimpled surface (right) obtained in the second generation of the reinforcement learning loop (Fig. 9). The color-maps for each dimple indicate the averaged temperature ( $T$ ) for their respective streamlines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

plausible alternatives within our parameter space. For instance, smaller dimples with high packing density can have more heat transfer area, yet their Nusselt number is inferior since their shapes do not enhance local heat transfer. Additionally, we have shown that at different Reynolds numbers, the optimal geometry is different as well. This showcases the

ability of our machine learning system to prioritize the alignment of roughness elements, and to select other appropriate surface features. Achieving similar results using traditional correlations would be difficult, since standard surface metrics (e.g. skewness) correspond to global averages, which are not sensitive to the exact location of roughness



**Fig. 13.** Instantaneous snapshot of the velocity field in the streamwise direction ( $U^+$ ) for the highest-performing dimpled surface of the second generation of the reinforcement learning study at  $Re_\tau = 180$ . The slice shown corresponds to the plane at  $z = L_z/2$ . Please note that  $U = U^+$  in our numerical framework. The bulk flow moves from left to right.



**Fig. 14.** Top ranked dimpled surfaces found during the optimization procedure for a friction Reynolds number of  $Re_\tau = 395$ . The label “previous best” refers to the top-ranked dimpled surface from the second generation of the reinforcement learning study for  $Re_\tau = 180$ . The bulk flow moves from left to right.






Fig. 15. Streamlines for the dimples with elliptical corners (top) shown in Fig. 14, and the top-ranked dimpled surface (bottom) obtained in the second generation of the reinforcement learning loop (Fig. 9) at  $Re_\tau = 180$ . The color-maps for each dimple indicate the averaged temperature ( $\bar{T}$ ) for their respective streamlines. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

elements. Thus, traditional correlations are not well-posed to predict how the alignment between dimples would enhance heat transfer.

In summary, we conclude that machine learning can be an effective tool to optimize rough surfaces for convective heat transfer enhancement. While typical reinforcement learning problems can require hundreds, if not thousands, of iterations to converge, our system can build highly effective surface designs in only a few iterations. Beyond the machine learning framework, our DNS analysis also shows that elongated dimples with opposite inclination angles are an interesting alternative to consider in optimization studies, since they produce a large enhancement in the average Nusselt number under equal pressure losses. Moreover, the local Nusselt numbers  $Nu(x, z)$  predicted for each configuration serve as an additional verification tool, to check whether the distributions follow physically plausible patterns, or if the mean Nusselt number  $\bar{Nu}$  are affected by spurious non-physical oscillations. Further design insights can also be obtained from  $Nu(x, z)$  to predict optimal heat transfer patterns, or to understand the sensitivity of the neural network to changes in the dimple geometry.

#### CRediT authorship contribution statement

**Rafael Diez Sanhueza:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jurriaan W.R. Peeters:** Writing – review & editing, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgments

This research did not receive additional support from an organization beyond the authors' academic institutions.

#### Appendix. Validation of the turbulent flow solver

In order to validate the new GPU-based DNS solver, a comparison was first performed with respect to the channel flow with a grit-blasted surface simulated by [27,28,32]. This DNS case uses a friction Reynolds number of  $Re_\tau = 180$  and a Prandtl number equal to unity. The grid size considered was  $(N_x \times N_y \times N_z) = (560 \times 280 \times 280)$  in each Cartesian direction, whereas the domain size is  $(L_x \times L_y \times L_z) = (5.63 \times 2 \times 2.815)$ . These settings are equivalent to a dimensionless grid size of  $\Delta x^+ = \Delta z^+ = 1.81$  in the streamwise and spanwise directions, respectively. In the wall normal direction, a constant value of  $\Delta y^+ = 0.65$  is kept near the walls. The validation data is plotted in Fig. A.16. Here, a comparison is performed with respect to the DNS data of Peeters and Sandham [27] and Busse et al. [28]. As it can be seen in the sub-plots, the new GPU-based DNS solver closely matches the turbulence statistics reported, and the associated changes in the velocity ( $\Delta U^+$ ) and temperature ( $\Delta T^+$ ) profiles.

To further validate the accuracy of our DNS solver, an additional DNS case was run replicating the work of Maaß and Schumann [45]. This case corresponds to a turbulent flow past a wavy surface. The results of the benchmark can be found in Fig. A.17. Here, it can be observed that the velocity profiles near the wavy surface are extremely close to the reference data. This is a strong validation of our implementation for the immersed boundary method, since the DNS case of Maaß and Schumann [45] operates at a relatively low Reynolds number with effects like flow separation.




Fig. A.16. Comparison between the turbulence statistics for the new GPU-based DNS solver and existing literature [27,28,32].

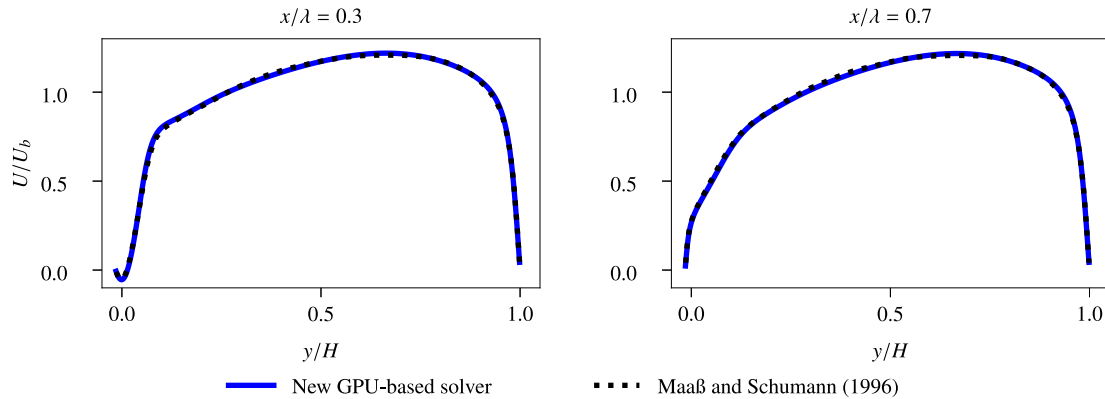


Fig. A.17. Validation of the GPU-accelerated solver with respect to the DNS case from Maaß and Schumann [45]. Here,  $H$  is the height of the channel and  $\lambda$  is the wavelength of the sinusoidal surface.

## Data availability

Data will be made available on request.

## References

- [1] D. Dipprey, R. Sabersky, Heat and momentum transfer in smooth and rough tubes at various prandtl numbers, *Int. J. Heat Mass Transfer* 6 (5) (1963) 329–353, URL: <https://www.sciencedirect.com/science/article/pii/0017931063900978>.
- [2] D. Chung, N. Hutchins, M.P. Schultz, K.A. Flack, Predicting the drag of rough surfaces, *Annu. Rev. Fluid Mech.* 53 (1) (2021) 439–471, [arXiv:https://doi.org/10.1146/annurev-fluid-062520-115127](https://doi.org/10.1146/annurev-fluid-062520-115127).
- [3] M.J. Walsh, L.M. Weinstein, Drag and heat-transfer characteristics of small longitudinally ribbed surfaces, *AIAA J.* 17 (7) (1979) 770–771, [arXiv:https://doi.org/10.2514/3.61216](https://doi.org/10.2514/3.61216).
- [4] P.M. Ligrani, M.M. Oliveira, T. Blaskovich, Comparison of heat transfer augmentation techniques, *AIAA J.* 41 (3) (2003) 337–362, [arXiv:https://doi.org/10.2514/2.1964](https://doi.org/10.2514/2.1964).
- [5] A. Rouhi, S. Endrikat, D. Modesti, R.D. Sandberg, T. Oda, K. Tanimoto, N. Hutchins, D. Chung, Riblet-generated flow mechanisms that lead to local breaking of Reynolds analogy, *J. Fluid Mech.* 951 (2022) A45.
- [6] P. Ligrani, Heat transfer augmentation technologies for internal cooling of turbine components of gas turbine engines, *Int. J. Rotating Mach.* 2013 (2013) 275653.
- [7] V. Afanasyev, Y. Chudnovsky, A. Leontiev, P. Roganov, Turbulent flow friction and heat transfer characteristics for spherical cavities on a flat plate, *Exp. Therm. Fluid Sci.* 7 (1) (1993) 1–8, URL: <https://www.sciencedirect.com/science/article/pii/089417779390075T>.
- [8] S.M. Choi, H.G. Kwon, H.M. Bae, H.K. Moon, H.H. Cho, Effects of staggered dimple array under different flow conditions for enhancing cooling performance of solar systems, *Appl. Energy* 342 (2023) 121120, URL: <https://www.sciencedirect.com/science/article/pii/S0306261923004841>.
- [9] Y. Rao, P. Zhang, Y. Xu, H. Ke, Experimental study and numerical analysis of heat transfer enhancement and turbulent flow over shallowly dimpled channel surfaces, *Int. J. Heat Mass Transfer* 160 (2020) 120195, URL: <https://www.sciencedirect.com/science/article/pii/S0017931020331318>.
- [10] A. Kumar, R. Maithani, A.R.S. Suri, Numerical and experimental investigation of enhancement of heat transfer in dimpled rib heat exchanger tube, *Heat Mass Transf.* 53 (12) (2017) 3501–3516.



- [11] Y. Rao, Y. Feng, B. Li, B. Weigand, Experimental and numerical study of heat transfer and flow friction in channels with dimples of different shapes, *J. Heat Transf.* 137 (3) (2015) 031901, [arXiv:https://asmedigitalcollection.asme.org/heattransfer/article-pdf/137/3/031901/6209299/ht\\_137\\_03\\_031901.pdf](https://asmedigitalcollection.asme.org/heattransfer/article-pdf/137/3/031901/6209299/ht_137_03_031901.pdf).
- [12] M. Jeong, M.Y. Ha, Y.G. Park, Numerical investigation of heat transfer enhancement in a dimpled cooling channel with different angles of the vortex generator, *Int. J. Heat Mass Transfer* 144 (2019) 118644, URL: <https://www.sciencedirect.com/science/article/pii/S0017931019311627>.
- [13] H. Xia, G. Tang, Y. Shi, W. Tao, Simulation of heat transfer enhancement by longitudinal vortex generators in dimple heat exchangers, *Energy* 74 (2014) 27–36, URL: <https://www.sciencedirect.com/science/article/pii/S0360544214002096>.
- [14] S. Xie, Z. Liang, L. Zhang, Y. Wang, A numerical study on heat transfer enhancement and flow structure in enhanced tube with cross ellipsoidal dimples, *Int. J. Heat Mass Transfer* 125 (C) (2018) 434–444.
- [15] T. Dagdevir, O. Keklikcioglu, V. Ozceyhan, Heat transfer performance and flow characteristic in enhanced tube with the trapezoidal dimples, *Int. Commun. Heat Mass Transfer* 108 (2019) 104299, URL: <https://www.sciencedirect.com/science/article/pii/S0735193319301654>.
- [16] K.-D.L. Abdus Samad, K.-Y. Kim, Shape optimization of a dimpled channel to enhance heat transfer using a weighted-average surrogate model, *Heat Transf. Eng.* 31 (13) (2010) 1114–1124, [arXiv:https://doi.org/10.1080/01457631003640453](https://doi.org/10.1080/01457631003640453).
- [17] Y. Rao, B. Li, Y. Feng, Heat transfer of turbulent flow over surfaces with spherical dimples and teardrop dimples, *Exp. Therm. Fluid Sci.* 61 (2015) 201–209, URL: <https://www.sciencedirect.com/science/article/pii/S0894177714002829>.
- [18] M.H. Cheraghi, M. Ameri, M. Shahabadi, Numerical study on the heat transfer enhancement and pressure drop inside deep dimpled tubes, *Int. J. Heat Mass Transfer* 147 (2020) 118845, URL: <https://www.sciencedirect.com/science/article/pii/S0017931019319519>.
- [19] M. Li, T.S. Khan, E. Al-Hajri, Z.H. Ayub, Single phase heat transfer and pressure drop analysis of a dimpled enhanced tube, *Appl. Therm. Eng.* 101 (2016) 38–46.
- [20] R.D. Sanhueza, I. Akkerman, J.W. Peeters, Machine learning for the prediction of the local skin friction factors and nusselt numbers in turbulent flows past rough surfaces, *Int. J. Heat Fluid Flow* 103 (2023) 109204, URL: <https://www.sciencedirect.com/science/article/pii/S0142727X23001030>.
- [21] R. Diez, Example of neural network to predict the local Nusselt number distributions in turbulent flows past rough surfaces, 2025, <https://github.com/Rafael10Diez/Machine-Learning-Distributed-Nusselt-Numbers-Efficient>.
- [22] A. Garai, S.M. Murman, Stabilization of the adjoint for turbulent flows, *AIAA J.* 59 (6) (2021) 2001–2013, [arXiv:https://doi.org/10.2514/1.J059998](https://doi.org/10.2514/1.J059998).
- [23] F. Salvadore, M. Bernardini, M. Botti, GPU accelerated flow solver for direct numerical simulation of turbulent flows, *J. Comput. Phys.* 235 (2013) 129–142, URL: <https://www.sciencedirect.com/science/article/pii/S0021999112006018>.
- [24] S. Isaev, A. Leontiev, M. Gritskovich, D. Nikushchenko, S. Guvernuyuk, A. Sudakov, K.-M. Chung, N. Tryaskin, M. Zubin, A. Sinyavin, Development of energy efficient structured plates with zigzag arrangement of multirow inclined oval trench dimples, *Int. J. Therm. Sci.* 184 (2023) 107988, URL: <https://www.sciencedirect.com/science/article/pii/S1290072922005166>.
- [25] S. Rashidi, F. Hormozi, B. Sundén, O. Mahian, Energy saving in thermal energy systems using dimpled surface technology – a review on mechanisms and applications, *Appl. Energy* 250 (2019) 1491–1547, URL: <https://www.sciencedirect.com/science/article/pii/S0306261919308268>.
- [26] P. Forrooghi, M. Stripf, B. Frohnapfel, A systematic study of turbulent heat transfer over rough walls, *Int. J. Heat Mass Transfer* 127 (2018) 1157–1168, URL: <https://www.sciencedirect.com/science/article/pii/S0017931018321148>.
- [27] J. Peeters, N. Sandham, Turbulent heat transfer in channels with irregular roughness, *Int. J. Heat Mass Transfer* 138 (2019) 454–467, URL: <https://www.sciencedirect.com/science/article/pii/S0017931018353456>.
- [28] A. Busse, M. Lützner, N.D. Sandham, Direct numerical simulation of turbulent flow over a rough surface based on a surface scan, *Comput. & Fluids* 116 (2015) 129–147, URL: <https://www.sciencedirect.com/science/article/pii/S0045793015001176>.
- [29] M. Thakkar, A. Busse, N. Sandham, Surface correlations of hydrodynamic drag for transitionally rough engineering surfaces, *J. Turbul.* 18 (2) (2017) 138–169, [arXiv:https://doi.org/10.1080/14685248.2016.1258119](https://doi.org/10.1080/14685248.2016.1258119).
- [30] P. Orlandi, S. Leonardi, DNS of turbulent channel flows with two- and three-dimensional roughness, *J. Turbul.* 7 (2006) N73, [arXiv:https://doi.org/10.1080/14685240600827526](https://doi.org/10.1080/14685240600827526).
- [31] J. Kim, P. Moin, R. Moser, Turbulence statistics in fully developed channel flow at low Reynolds number, *J. Fluid Mech.* 177 (1987) 133–166.
- [32] M. Thakkar, A. Busse, N. Sandham, Direct numerical simulation of turbulent channel flow over a surrogate for Nikuradse-type roughness, *J. Fluid Mech.* 837 (2018).
- [33] S. Leonardi, P. Orlandi, L. Djenidi, R.A. Antonia, Heat transfer in a turbulent channel flow with square bars or circular rods on one wall, *J. Fluid Mech.* 776 (2015) 512–530.
- [34] E. Fadlun, R. Verzicco, P. Orlandi, J. Mohd-Yusof, Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations, *J. Comput. Phys.* 161 (1) (2000) 35–60, URL: <https://www.sciencedirect.com/science/article/pii/S0021999100964842>.
- [35] R.D. Sanhueza, J. Peeters, P. Costa, A pencil-distributed finite-difference solver for extreme-scale calculations of turbulent wall flows at high Reynolds number, 2025, URL: <https://arxiv.org/abs/2502.06296>.
- [36] M. Yang, J.-H. Kang, K.-H. Kim, O.-K. Kwon, J.-I. Choi, PaScaLTDMA 2.0: A multi-GPU-based library for solving massive tridiagonal systems, *Comput. Phys. Comm.* 290 (2023) 108785, URL: <https://www.sciencedirect.com/science/article/pii/S0010465523001248>.
- [37] K.-H. Kim, J.-H. Kang, X. Pan, J.-I. Choi, PaScaL-TCS: A versatile solver for large-scale turbulent convective heat transfer problems with temperature-dependent fluid properties, *Comput. Phys. Comm.* 290 (2023) 108779, URL: <https://www.sciencedirect.com/science/article/pii/S0010465523001248>.
- [38] P. Costa, CaNS Version 3.0: A code for fast, massively-parallel direct numerical simulations (DNS) of canonical flows, 2025, <https://github.com/CaNS-World/CaNS/releases/tag/v3.0.0>.
- [39] S. Isaev, A. Leontiev, Y. Chudnovsky, D. Nikushchenko, I. Popov, A. Sudakov, Simulation of vortex heat transfer enhancement in the turbulent water flow in the narrow plane-parallel channel with an inclined oval-trench dimple of fixed depth and spot area, *Energies* 12 (7) (2019) URL: <https://www.mdpi.com/1996-1073/12/7/1296>.
- [40] Y. Chen, Y.T. Chew, B.C. Khoo, Enhancement of heat transfer in turbulent channel flow over dimpled surface, *Int. J. Heat Mass Transfer* 55 (25–26) (2012) 8100–8121.
- [41] M.A. Elyan, A. Rozati, D.K. Tafti, Investigation of dimpled fins for heat transfer enhancement in compact heat exchangers, *Int. J. Heat Mass Transfer* 51 (11) (2008) 2950–2966, URL: <https://www.sciencedirect.com/science/article/pii/S0017931007005893>.
- [42] J. Turnow, N. Kornev, V. Zhdanov, E. Hassel, Flow structures and heat transfer on dimples in a staggered arrangement, *Int. J. Heat Fluid Flow* 35 (2012) 168–175.
- [43] K.S.Y. Zhengyi Wang, B.C. Khoo, DNS of low Reynolds number turbulent flows in dimpled channels, *J. Turbul.* 7 (2006) N37, [arXiv:https://doi.org/10.1080/14685240600595735](https://doi.org/10.1080/14685240600595735).
- [44] J. Bons, A critical assessment of Reynolds analogy for turbine flows, *J. Heat Transf.* 127 (5) (2005) 472–485, [arXiv:https://asmedigitalcollection.asme.org/heattransfer/article-pdf/127/5/472/5734392/472\\_1.pdf](https://asmedigitalcollection.asme.org/heattransfer/article-pdf/127/5/472/5734392/472_1.pdf).
- [45] C. Maaß, U. Schumann, Direct Numerical Simulation of Separated Turbulent Flow over a Wavy Boundary, Vieweg+Teubner Verlag, Wiesbaden, 1996, pp. 227–241, URL: [https://doi.org/10.1007/978-3-322-89849-4\\_17](https://doi.org/10.1007/978-3-322-89849-4_17).