Progress in Nuclear Energy 191 (2026) 106100

Contents lists available at ScienceDirect

Progress in Nuclear Energy

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

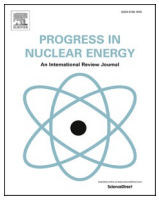

Explainable critical heat flux prediction in vertical circular channels by self-attention-enhanced one-dimensional convolutional neural networks

Ibrahim Ahmed [a,*], Enrico Zio [a,b]

[a] Department of Energy, Politecnico di Milano, Milan, Italy
[b] MINES Paris-PSL University, CRC, Sophia Antipolis, France

ARTICLE INFO

Keywords:
Critical heat flux
Thermohydraulic systems
Nuclear reactors
Explainability
Interpretability
Attention mechanisms
Convolutional neural networks

ABSTRACT

Critical heat flux (CHF) is an important parameter in thermohydraulic systems for both nuclear and industrial applications. Given the importance of CHF, explainability in its prediction becomes necessary for those safety-critical applications like water-cooled nuclear reactors. For a predictive model to be deemed explainable, it must not only make accurate predictions but also elucidate the rationale behind specific predictions. In this paper, we propose an explainable self-attention-enhanced one-dimensional convolutional neural network (SA-1DCNN) for CHF prediction in vertical circular channels. The proposed hybrid architecture integrates convolutional layers for extracting localized patterns with self-attention mechanisms for capturing global contextual dependencies among key physical parameters. A Bayesian Optimization algorithm is used to automatically set the hyperparameters of the proposed SA-1DCNN model, resulting in an optimized, robust and generalized architecture for CHF prediction. Beyond good predictive performance, explainability and interpretability are also achieved through comprehensive analyses, including mean activation profiling, spatial attention mapping and layer-wise correlation analysis. The proposed SA-1DCNN model is validated on the experimental CHF data provided by the Working Party on Scientific Issues and Uncertainty Analysis of Reactor Systems (WPRS) Expert Group on Reactor Systems Multi-Physics (EGMUP) task force on AI and ML for Scientific Computing in Nuclear Engineering projects, supported by the OECD/NEA. The obtained results show that the proposed SA-1DCNN model not only outperforms state-of-the-art methods in predictive performance but also offers transparent and physically consistent insights into the mechanistic and thermohydraulic phenomena governing CHF. This combination of predictive power, explainability and interpretability establishes the proposed model as a promising tool for data-driven decision-making in nuclear thermohydraulic systems.

1. Introduction

Critical Heat Flux (CHF) is the upper thermal margin in boiling heat transfer processes, beyond which a severe degradation in heat transfer efficiency occurs (Todreas and Kazimi, 2011). This phenomenon is commonly referred to as “boiling crisis” and defines the conversion from nucleate boiling to film boiling (Bruder et al., 2017), leading to overheating and unanticipated phase changes. Consequently, the CHF constitutes a fundamental safety parameter in critical thermohydraulic systems like water-cooled nuclear reactors, heat exchange devices, boilers and steam generators (Chang and Baek, 2003).

Especially for water-cooled nuclear power plants (NPPs), ensuring safe and efficient operation requires careful consideration of thermohydraulic safety margins like those defined with reference to CHF (Chang and Baek, 2003), and their regulation within stringent operational limits. With regards to CHF, the *Departure from Nucleate Boiling Ratio* (DNBR), defined as the ratio of predicted CHF to the operating local heat flux, is widely used for this purpose. To ensure safe operating margins, the *Minimum DNBR* (MDNBR) must exceed unity. For example, the U.S. Nuclear Regulatory Commission (NRC) requires at least 1.3 value of MDNBR when employing the W-3 correlation for CHF prediction (Todreas and Kazimi, 2011). Although there is a general consensus that the flow characteristics near the heated surface predominantly influence the CHF initiation (Bruder et al., 2017), the complexity of the physical mechanisms responsible for it has impeded the development of a unified predictive model. Then, CHF prediction is typically obtained by various empirical/semi-empirical correlations (Bruder et al., 2017). The earliest prediction approaches comprised empirical correlations

* Corresponding author.
E-mail addresses: ibrahim.ahmed@polimi.it (I. Ahmed), enrico.zio@polimi.it, enrico.zio@minesparis.psl.eu (E. Zio).

https://doi.org/10.1016/j.pnucene.2025.106100
Received 17 July 2025; Received in revised form 8 September 2025; Accepted 22 October 2025
Available online 28 October 2025

(Biasi et al., 1968; Katto, 1992; Tong, 1967) and mechanistic models (Celata et al., 1999; Lee and Mudawwar, 1988; Weisman and Pei, 1983). Empirical correlations, which are derived from experimental datasets, are relatively straightforward and effective under conditions closely aligned with their calibration domains. However, their applicability in practice is often constrained by their reliance on empirically derived parameters. Mechanistic models, instead, strive to incorporate fundamental physical phenomena such as flow regimes and bubble dynamics. Nevertheless, due to an incomplete understanding of CHF mechanisms, these models frequently lack the accuracy necessary for practical applications and must depend on additional empirical inputs. Given the limitations of analytical and empirical/semi-empirical models, particularly under transient conditions or in rod bundles with non-uniform heat flux, empirical approaches like Look-Up Tables (LUTs) have received significant attraction. The initial LUT development by Doroshchuk in 1975 (Doroshchuk et al., 1975) was based on 5000 data points and has since then been continually refined with new experimental data. Notice that Groeneveld et al. produced the most widely adopted LUT in 2006 (Groeneveld, 2019; Groeneveld et al., 2007), which maps CHF as a function of mass flux, pressure and outlet quality. Despite its broad coverage and ease of use, the LUT approach lacks physical interpretability and its accuracy suffers in data-scarce regions. Nevertheless, this 2006 LUT remains a standard reference in nuclear thermohydraulic systems (Groeneveld, 2019).

With recent advancements in artificial intelligence (AI) and machine learning (ML) techniques, innovative solutions for solving complex problems in predictive modelling have been developed. In (Mazzola, 1997), artificial neural networks (ANNs) were employed for the estimation of the parameters of empirical correlation, which are typically derived using best-fit methods. This modelling approach maintained the benefits of analytical analysis while partially alleviating the black-box nature of direct use of ANNs, which were generally confined to tuning of empirical correlation. In (Lee et al., 2006), a backpropagation neural network (BPN) was utilized for the prediction of the MDNBR. Initially, a thermohydraulic code was used for the simulation of a Total Loss of Flow Accident (TLOFA), a critical condition resulting in the most severe MDNBR. The results from the TH code were, then, utilized to provides in inputs to an auto-associative neural network (AANN) to refine their values. Specifically, the BPN used these refined values for the prediction of the MDNBR, which demonstrates high accuracy. In (Jiang and Zhao, 2013), a hybrid model that combined support vector regression (ν-SVR) with radial basis function networks (RBFNs) was developed for CHF prediction, and the results demonstrated improved performance when compared to standard SVR and empirical correlation methods. In (He and Lee, 2018), the proposed method in (Jiang and Zhao, 2013) was extended to address sparsely distributed data, which achieved enhanced accuracy by concentrating training near critical inflection points. In (Khalid et al., 2023), the performances of random forest (RF), ANNs and support vector machines (SVMs) were evaluated. The results showed that the combined used of ANNs and LUTs (hybrid model) outperformed other models. In (Zhao et al., 2020), a physics-informed ML (PIML) was proposed. The proposed approach improved both generalization abilities and accuracy when compared to traditional models. In (Mao and Jin, 2024), ML models were integrated with physics-based approaches to develop a PIML framework for predicting CHF. The work demonstrated that neural networks (NNs) augmented with LUT information provided superior stability and robustness compared to other models. In (Khalid et al., 2024), a novel ensemble approach was proposed, wherein deep sparse autoencoders (AEs) were employed extracting features from input data and subsequently coupled with a deep neural network (DNN) serving as a meta-learner for predicting CHF. This methodology effectively addressed the issues of state-of-the-art models by utilizing a large and diverse dataset containing a broad regime of operating conditions, thereby substantially improving the predictive accuracy and generalizability of the model. A detailed review of AI/ML methods for CHF prediction is provided in (Qi et al., 2025).

Recently, the OECD Nuclear Energy Agency (OECD/NEA) has launched a dedicated task force on AI and ML for Scientific Computing in Nuclear Engineering (OECD/NEA, 2022), which has made available benchmark dataset for CHF prediction (Le Corre et al., 2024). The primary objectives of the task force include (OECD/NEA, 2022): *1*) to build benchmarks for performance evaluation of AI/ML methods in multi-physics Modeling and Simulation (M&S) of reactor systems, and *2*) to provide strategic recommendations to the Working Party on Scientific Issues and Uncertainty Analysis of Reactor Systems (WPRS) and the broader nuclear community regarding the scientific development needed (methods, benchmark exercises and data) to enable trustworthy AI/ML applications in nuclear scientific computing challenges. This benchmark dataset has already been used in recent research for CHF prediction to demonstrate its practical relevance. In (Zhou et al., 2024), advanced AI techniques, such as ANNs, Transformer and convolutional neural networks (CNNs) models, were applied to the prediction of CHF. Among these, the Transformer-based model, which leverages self-attention mechanisms to dynamically assign relevance to different segments of the input data, demonstrated superior predictive performance. In (Ahmed et al., 2025b), an optimized ensemble of NNs was proposed to improve CHF prediction by addressing challenges related to ensemble model selection and complexity reduction. The authors introduced two systematic procedures for identifying high-performing candidate NNs and combining them into an optimal ensemble. Validation using a CHF benchmark dataset showed that the ensemble approach outperformed individual models and other leading AI-based methods in terms of prediction accuracy. In (Wang et al., 2025), three deep learning architectures, including Transformer, Mamba and Temporal Convolutional Network (TCN), were evaluated for CHF prediction. The study identified Transformers as the most accurate CHF prediction model. The analysis on the input parameter selection through mechanistic analysis significantly enhanced prediction accuracy.

While recent advancements in AI/ML have significantly improved the prediction of CHF compared to traditional approaches, limited research has addressed the explainability (i.e., the ability to communicate algorithmic decisions and their data-driven basis to end users) and interpretability (i.e., the capability of the model to adhere to domain-specific structural knowledge, including monotonicity, causality and physical constraints) of these models. Although studies such as (Myers et al., 2025; Sajjad et al., 2025) have investigated explainable AI (XAI) for nuclear power applications, the efforts have primarily been focused on input-output relationships (i.e., identifying which input variables most influence the output) without probing the internal representations or mechanisms by which these models operate. To bridge this gap, we propose an explainable self-attention-enhanced one-dimensional convolutional neural network (SA-1DCNN) for CHF prediction. This architecture not only makes accurate predictions, but it also provides us with insights on the internal operations of the model, which makes it more trustworthy and transparent in safety-critical applications. The proposed hybrid architecture integrates convolutional layers for extracting localized patterns with self-attention mechanisms (Vaswani et al., 2017) for capturing global contextual dependencies among key thermohydraulic parameters. One-dimensional CNNs (1DCNNs) are suitable for this task, as they are effective in hierarchically extracting features from sequential or one-dimensional inputs (Kiranyaz et al., 2021), such as time-series data or one-dimensional thermodynamic parameter profiles. The incorporation of a self-attention mechanism in the proposed model introduces several advantages for CHF modeling.

- it enables the model to simultaneously capture local patterns and global interdependencies among input parameters, enhancing predictive performance under complex two-phase flow conditions;
- the attention mechanism dynamically assigns contextual relevance to input features, thereby improving model interpretability;

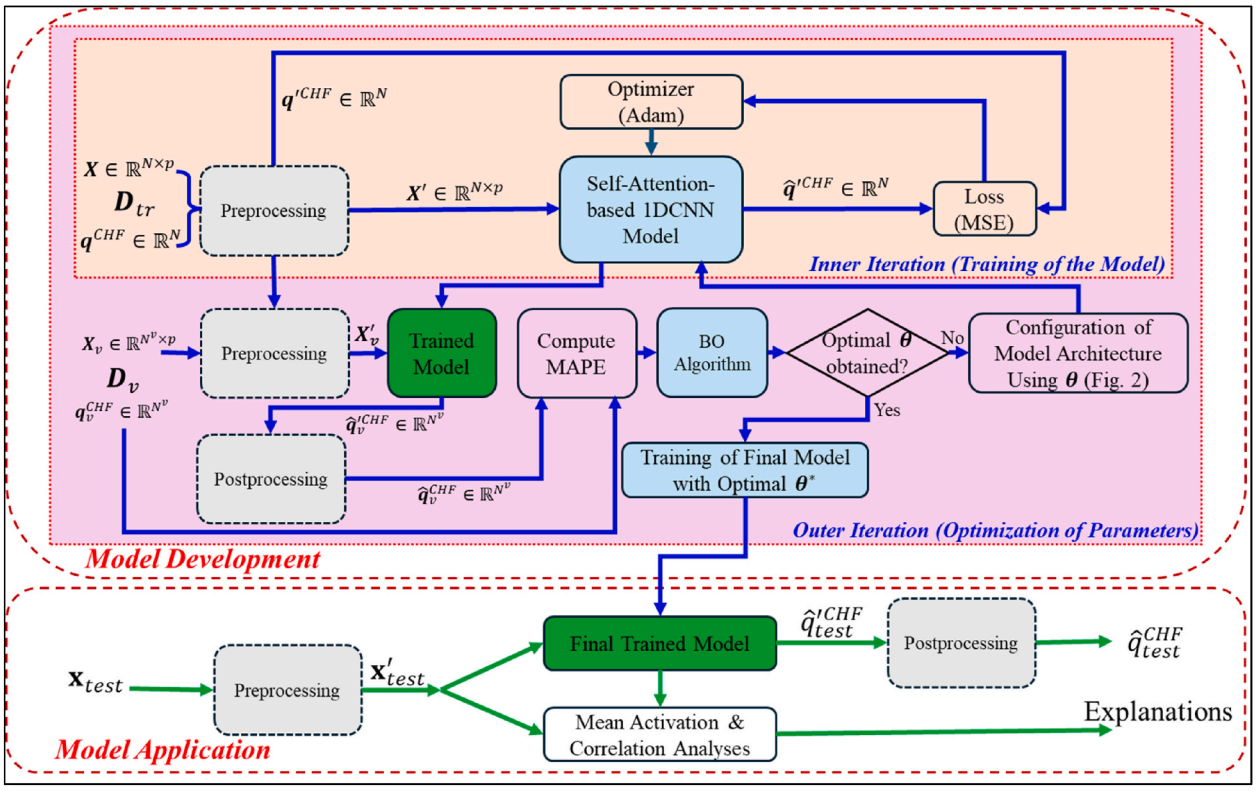

Fig. 1. Flow diagram of the proposed explainable SA-1DCNN method for CHF prediction.

- the hybrid architecture (SA-1DCNN) integrates convolutional feature extraction with attention-based global reasoning, promoting generalization across diverse operational regimes;
- the model efficiently scales with input complexity and varying sequence lengths, making it suitable for large experimental datasets;
- the attention scores serve as interpretable indicators of feature importance, thereby supporting physics-informed insights and facilitating transparent deployment in nuclear safety-critical applications.

In addition, Bayesian Optimization (BO) technique is used to automatically set the hyperparameters of the proposed SA-1DCNN model, resulting in an optimized, robust and generalized architecture for CHF prediction. Importantly, explainability and interpretability are achieved through comprehensive analyses, including mean activation profiling, spatial attention mapping and layer-wise correlation analysis. Compared to other DL models, the proposed SA-1DCNN model offers significantly enhanced transparency. This facilitates a deeper understanding of the physical phenomena that govern CHF and supports trustworthy, data-driven decisions in the designs and operations of thermohydraulic systems.

The proposed SA-1DCNN model is validated on experimental CHF data originally presented in (Groeneveld, 2019) and provided by the WPRS Expert Group on Reactor Systems Multi-Physics (EGMUP) task force on AI and ML for Scientific Computing in Nuclear Engineering project, a project promoted by the OECD/NEA (Le Corre et al., 2024). The obtained results indicate that the proposed SA-1DCNN method not only outperforms state-of-the-art modelling approaches in terms of predictive performance, but also offers transparent and physically consistent information about the thermohydraulic mechanisms that govern CHF. This combination of predictive power, explainability and interpretability establishes the proposed method as a promising tool for data-driven decision-making in nuclear thermohydraulic systems.

The remainder of this paper is organized as follows. In Section 2, the problem statement is presented. In Section 3, the proposed SA-1DCNN method is described. The application case study is described in Section 4. The results and discussion are given in Section 5. Finally, Section 6 summarizes the concluding remarks on the performed work.

2. Problem statement

In the design phase of a water-cooled nuclear reactor, it is important to accurately predict the CHF to ensure efficient heat transfer and to guarantee reliable safety margins. This is done with reference to a set of coherent experimental measurements $\mathbf{x}$, which depend on P relevant physical parameters under specified boundary conditions at a given observation i:

$$\mathbf{x}_i = \left[x_{i,1}, \quad \ldots, \quad x_{i,j}, \quad \ldots, \quad x_{i,P} \right] \tag{1}$$

These measured parameters comprise both hydraulic and geometrical parameters, specifically including the pressure, tube diameter, mass flux, heated length and outlet quality. These are the parameters that influence the CHF behavior, indicated as q_i^{CHF}, at the same observation i.

The dataset $\boldsymbol{D}_{tr} = \left[\mathbf{x}_i \quad q_i^{CHF} \right]_{i=1,\ldots,N}$ is available and assumed to contain the experimentally measured variables collected over a specific period of time, comprising:

- the experimental observations $\boldsymbol{X} \in \mathbb{R}^{N \times P}$, whose element $x_{i,j}$ ($i = 1, \ldots, N$ and $j = 1, \ldots, P$) is the measured parameter j at the observation i.
- the corresponding CHF values $\boldsymbol{q}^{CHF} \in \mathbb{R}^N$, containing the measured CHF value q_i^{CHF} ($i = 1, \ldots, N$) at each observation i.

Consider a new input $\mathbf{x}_{test}$ measured at the present observation time, the objective of this work is to develop a predictive model that takes as input $\mathbf{x}_{test}$ and predicts as output the corresponding CHF value $\hat{q}_{test}^{CHF}$.

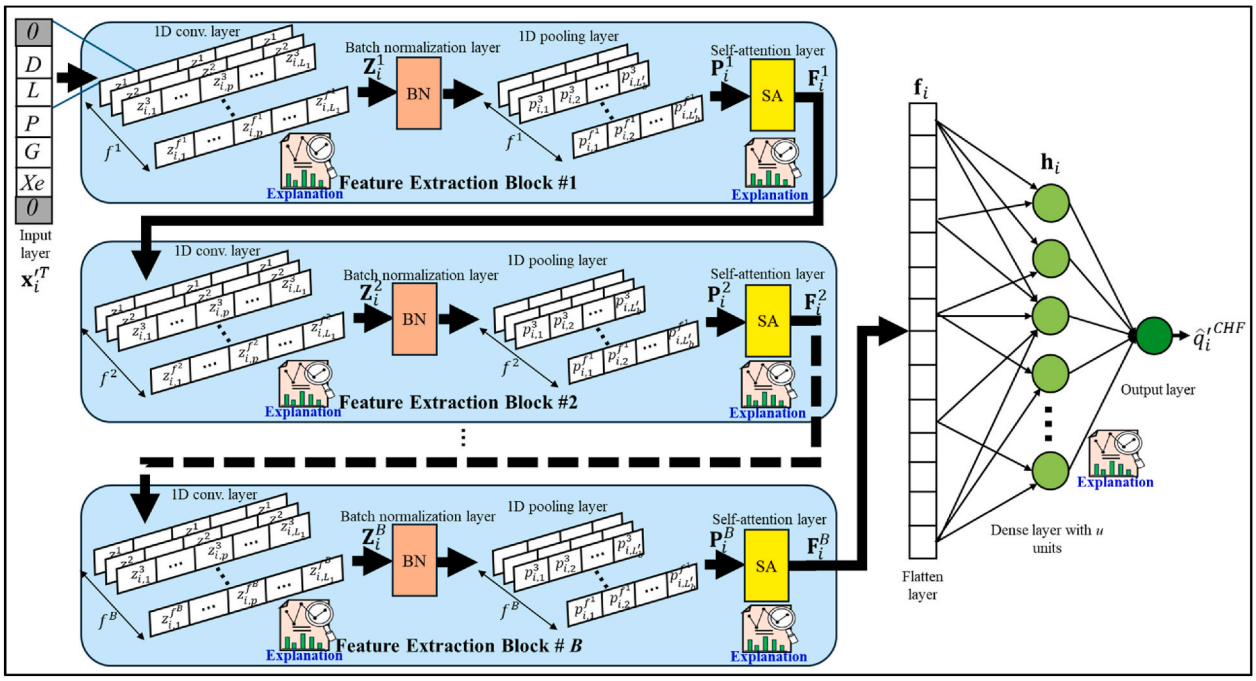

Fig. 2. Block diagram of the proposed self-attention-enhanced 1DCNN architecture.

3. Proposed method

The proposed explainable CHF prediction method is sketched in Fig. 1. It comprises the development of self-attention-enhanced one-dimensional convolutional neural networks (SA-1DCNN) model (Section 3.1), and the application and explanation of the developed model (Section 3.2).

3.1. Development of a self-attention-enhanced 1DCNN model for CHF prediction

3.1.1. Preprocessing of model input and output

To ensure numerical stability and facilitate effective training of the model, both input and output parameters are appropriately transformed. All input variables, including geometric and thermohydraulic parameters, are normalized to the range $[-1, 1]$ using min-max scaling:

$$x'_{i,j} = 2 \cdot \frac{x_{i,j} - x_j^{min}}{x_j^{max} - x_j^{min}} - 1 \tag{2}$$

where x_j^{min} and x_j^{max} are the minimum and maximum values of the input parameter j in the training dataset. This transformation centers the data around zero and scales all parameters to lie within the same bounded interval, which is particularly important for NN-based models where the activation functions are sensitive to input scale (LeCun et al., 2012). Unlike standard normalization approaches, min-max scaling preserves the original distribution shape and relative spacing of the data, making it preferable for methods where assumptions of Gaussianity are not required (Han et al., 2012). Furthermore, scaling to a symmetric interval such as $[-1, 1]$ has been shown to facilitate faster convergence and reduce the likelihood of vanishing or exploding gradients in deep learning models (Lecun et al., 2015).

Given the wide dynamic range and skewed distribution of the output variable (CHF), a logarithmic transformation is used to stabilize the variance and improve the learning process of the proposed method:

$$q_i'^{CHF} = \ln\left(q_i^{CHF} + \epsilon\right) \tag{3}$$

where ϵ is a small constant value set equal to $\epsilon = 10^{-6}$ to avoid undefined logarithmic values near-zero. This transformation reduces heteroscedasticity (Osborne, 2010) and improves the generalization ability of the proposed model across diverse operating conditions.

3.1.2. Architecture of the SA-1DCNN model

To predict the CHF from five physical input parameters: tube diameter (D), pressure (P), mass flux (G), heated length (L) and outlet quality (Xe), we develop a 1DCNN architecture enhanced with self-attention mechanisms (SA-1DCNN). This architecture is designed to extract both localized (via convolution) and global (via attention) contextual features from the input sequence and to generate accurate CHF predictions in a data-efficient and interpretable manner. The architecture of the proposed SA-1DCNN model is detailed in Fig. 2. The model comprises B sequential blocks, each consisting of a 1D convolutional layer, batch-normalization, max-pooling and a self-attention mechanism. These are followed by a fully-connected layer for CHF prediction.

Mathematically, the proposed SA-1DCNN model is described as follows. First, each normalized input vector, $\mathbf{x}_i'^T \in \mathbb{R}^{5\times 1}$ is treated as a univariate sequence with one feature channel processed through B sequential feature extraction blocks. For each block b ($b = 1, 2, \cdots, B$), the output of the 1D convolutional layer is computed as:

$$\mathbf{Z}_i^b = ReLU\left(Conv1D\left(\mathbf{F}_i^{b-1};\ \mathbf{W}_c^b;\ \mathbf{b}_c^b\right)\right) \in \mathbb{R}^{L_b \times f^b} \tag{4}$$

where $Conv1D(\cdot)$ is the 1D convolutional operator; $ReLU(\cdot)$ is the Rectified Linear Unit (ReLU) activation function; $\mathbf{F}_i^{b-1}$ is the output activations (feature maps) of the previous block $b-1$ (or $\mathbf{x}_i'^T$ for $b = 1$); $\mathbf{W}_c^b \in \mathbb{R}^{k\times 1\times f^b}$ and $\mathbf{b}_c^b \in \mathbb{R}^{f^b}$ are the convolutional weights and biases (with k and f^b being the filter size and number of filters, respectively); L_b is the sequence length after convolution, depending on padding and stride used. This output is passed through batch normalization (BN) to enhance training stability and max-pooling to reduce dimensionality along the sequence length by taking maximum values over non-overlapping regions:

$$\mathbf{P}_i^b = MaxPool1D\left(BN\left(\mathbf{Z}_i^b\right)\right) \in \mathbb{R}^{L_b' \times f^b} \tag{5}$$

To capture contextual dependencies and interactions among features, a self-attention (SA) mechanism is applied. Self-attention is a mechanism used in ML to capture dependencies and relationships within input sequences (Vaswani et al., 2017). It enables the proposed model to identify and weigh the relevance of different parts of the input parameters by attending to itself. Specifically, the SA mechanism computes the attention output as:

$$\mathbf{F}_i^b(\mathbf{P}_i^b) = softmax\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}\right)\mathbf{V} \in \mathbb{R}^{L_b' \times f^b} \tag{6}$$

where $\mathbf{Q}$, $\mathbf{K}$ and $\mathbf{V}$ are the query, key and value matrices, and $d = f^b$ is the attention dimension. Here, the matrices $\mathbf{Q}$, $\mathbf{K}$ and $\mathbf{V}$ are set to the output of the previous layer (i.e., $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \mathbf{P}_i^b$), resulting in a self-attention configuration. This configuration allows the model to learn the internal dependencies within the local feature sequence, enhancing its capacity to capture non-local patterns and interactions that may not be well represented by convolutional operations alone. The output, $\mathbf{F}_i^B$ from the final block B, is flattened into a vector:

$$\mathbf{f}_i = Flatten\left(\mathbf{F}_i^B\right) \in \mathbb{R}^{L_b' f^b} \tag{7}$$

and a fully-connected layer transforms this high-dimensional flattened vector into a lower-dimensional representations:

$$\mathbf{h}_i = ReLU\left(\mathbf{W}_{fc}\mathbf{f}_i + \mathbf{b}_{fc}\right) \in \mathbb{R}^u \tag{8}$$

where $\mathbf{W}_{fc} \in \mathbb{R}^{u \times (L_b' f^b)}$ and $\mathbf{b}_{fc} \in \mathbb{R}^u$ are the weight matrix and bias vector (with u being the number of hidden nodes in the layer). Finally, the output layer maps this lower-dimensional representation, $\mathbf{h}_i$ to a CHF prediction:

$$\widehat{q}_i^{\,'CHF} = \phi\left(\mathbf{w}_{out}^T\mathbf{h}_i + b_{out}\right) \in \mathbb{R}^1 \tag{9}$$

where $\mathbf{w}_{out} \in \mathbb{R}^u$ and $b_{out} \in \mathbb{R}^1$ are the weight and bias vectors of the output layer, and $\phi(\cdot)$ is the output activation function.

To ensure that the model predicts strictly positive CHF values and remains consistent with the scaling transformation employed in Eq. (3), a parameterized Softplus activation function is applied for the output layer:

$$\phi_{\beta,\alpha}(x) = \beta \cdot \ln(1 + \exp(x)) + \alpha \tag{10}$$

where β and α are the scaling and offset (bias constant) hyperparameters that can be adjusted during the training process. This proposed variation of the softplus activation function allows adjusting the smoothness and position of the activation function, thereby adapting to the specific output ranges.

The trainable model parameters $\mathbf{W} = \left\{\mathbf{W}_c^b;\ \mathbf{W}_{fc};\ \mathbf{w}_{out}^T\right\}$ and $\mathbf{b} = \left\{\mathbf{b}_c^b;\ \mathbf{b}_{fc};\ b_{out}\right\}$ of the above described SA-1DCNN model are set by employing the error backpropagation technique (Rumelhart et al., 1986) using Adam optimizer (inner iteration of Fig. 1), which minimizes the mean squared error (MSE) used as loss function:

$$\mathscr{L}_{1DCNN}(\mathbf{W}, \mathbf{b}) = MSE = \frac{1}{N}\sum_{i=1}^{N}\left(q_i^{'CHF} - \widehat{q}_i^{\,'CHF}\right)^2 \tag{11}$$

where $q_i^{'CHF}$ and $\widehat{q}_i^{\,'CHF}$ are the normalized measured and predicted CHF values, respectively.

The transformed target, $q_i^{'CHF}$ is used during the training process and the model final predictions are subsequently exponentiated (post-processing) to recover the original scale of the CHF:

$$\widehat{q}_i^{CHF} = \exp\left(\widehat{q}_i^{\,'CHF}\right) - \epsilon \tag{12}$$

3.1.3. Optimization of the SA-1DCNN hyperparameters

To achieve optimal predictive performance of the proposed model, we employ Bayesian Optimization (BO) to effectively search the hyperparameter space. BO is a probabilistic optimization framework that efficiently balances exploration and exploitation by constructing a surrogate model (typically a Gaussian Process (GP)) over the objective function and selecting optimal configurations through an acquisition function (Močkus, 1974). In this work, the tunable hyperparameters, $\boldsymbol{\theta}$ to be optimized include:

- Number of feature extraction blocks B;
- Number of filters per block $f^b (b = 1, 2, \cdots, B)$, determining the representational capacity of each convolutional layer;
- Convolutional kernel (filter) size k common for all blocks, affecting the local receptive field for spatial feature extraction;
- Number of hidden nodes in the fully-connected layer u, controlling the ability of the network to capture global non-linear relationships;
- Learning rate η, governing the speed and stability of convergence during backpropagation;
- Batch size m, influencing the gradient estimation quality and computational efficiency;
- Output activation function parameters (Eq. (10)): β, which controls the curvature and smoothness of the function and α, which adjusts the vertical shift.

BO is chosen over traditional methods such as grid search and random search due to its sample efficiency and adaptive learning capability, especially valuable in this case where model evaluations are computationally expensive and the objective landscape is non-convex and noisy (Wang et al., 2020). Furthermore, unlike evolutionary algorithms, BO requires fewer iterations to find near-optimal solutions (Ahmed et al., 2025a), and also enables uncertainty quantification (Wang et al., 2020). BO is here applied to iteratively search for the best combination of the above described hyperparameters $\boldsymbol{\theta}$ to minimize the objective function Mean Absolute Percentage Error (MAPE) on the validation set $\boldsymbol{D}_v = \left[\mathbf{x}_i \quad q_i^{CHF}\right]_{i=1,\ldots,N^v}$ (outer iteration of Fig. 1):

$$\mathscr{L}_{BO}(\boldsymbol{\theta}) = MAPE = \frac{100}{N^v}\sum_{i=1}^{N^v}\left|\frac{\widehat{q}_i^{CHF} - q_i^{CHF}}{q_i^{CHF}}\right| \tag{13}$$

MAPE is suitable for CHF due to its wide dynamic range of interpretability in terms of relative prediction error. In this approach, the model is wrapped as a black-box function that maps a hyperparameter configuration to a validation error score. BO maintains a surrogate model (GP) of this function to estimate the posterior distribution of the objective function. An acquisition function, such as Upper Confidence Bound (UCB), Expected Improvement (EI), Probability of Improvement (PI) and Thompson Sampling (TS), guides the selection of the next candidate configuration by balancing exploration and exploitation. The proposed hyperparameter optimization based on BO is summarized in Algorithm 1. In each iteration k (outer iteration), SA-1DCNN model architecture is configured with the given hyperparameters $\boldsymbol{\theta}_k$, trained on the training set by minimizing Eq. (11) (inner iteration) and evaluated on the validation set to compute MAPE (the objective function). Finally, the obtained optimal hyperparameters are used to develop and train the final model for CHF prediction.

Algorithm 1. SA-1DCNN Hyperparameters setting by BO

Inputs:
$\mathcal{H}$ /*Hyperparameter search space*/
$\mathcal{L}_{BO}(\boldsymbol{\theta})$ /*Objective function: MAPE on validation set*/
T /*Number of iterations*/;
N_0 /*Initial samples (random or Latin hypercube)*/

Output: $\boldsymbol{\theta}^* \in \mathcal{H}$ /*Optimal hyperparameters*/

Begin
  Initialize the dataset $\boldsymbol{D} = \{\boldsymbol{\theta}_i, \mathcal{L}_{BO}(\boldsymbol{\theta}_i)\}_{i=1}^{N_0}$ using N_0 samples
  Fit a surrogate model $\mathcal{S}$ (e.g., GP) to $\boldsymbol{D}$: $\mathcal{S} \leftarrow \mathrm{Fit}(\boldsymbol{D})$
  For $k = 1$ to T **do** /**iteration loop (Outer iteration)*/
    Select the next candidate: $\boldsymbol{\theta}_k = \arg\max_{\boldsymbol{\theta}\in\mathcal{H}}\big(\mathrm{AquisitionFunction}(\boldsymbol{\theta};\ \mathcal{S})\big)$
    Configure the model (Figure 2) with configuration $\boldsymbol{\theta}_k$ /*Inner iteration*/
    Train the configured architecture by minimizing Eq. (11) /*Inner iteration*/
    Evaluate the model on the validation set $\boldsymbol{D}_v$ to compute $\mathcal{L}_{BO}(\boldsymbol{\theta}_k)$ using Eq. (13)
    Augment $\boldsymbol{D} \leftarrow \boldsymbol{D} \cup \{\boldsymbol{\theta}_k, \mathcal{L}_{BO}(\boldsymbol{\theta}_k)\}$
    Update surrogate model $\mathcal{S} \leftarrow \mathrm{Fit}(\boldsymbol{D})$
  End
  Return the best configuration of hyperparameters: $\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}\in\boldsymbol{D}}\big(\mathcal{L}_{BO}(\boldsymbol{\theta})\big)$
End

3.2. Application and explanation of the model

3.2.1. Prediction of the CHF by the proposed method

Given a new test measurements of the physical quantities $\mathbf{x}_{test}$ taken at the current observation, the proposed method (Fig. 1, bottom) first preprocesses it to $\mathbf{x}'_{test}$ using Eq. (2); then, the transformed vector $\mathbf{x}'_{test}$ is fed into the trained SA-1DCNN model for the prediction of the CHF $\widehat{q}'^{CHF}_{test}$; subsequently, the prediction $\widehat{q}'^{CHF}_{test}$ is postprocessed to obtain the original scale of the CHF, $\widehat{q}^{CHF}_{test}$ using Eq. (12).

3.2.2. Explainability analysis

To enhance and ensure scientific transparency and foster the needed model trustworthiness, we analyze the intermediate features learned by the proposed SA-1DCNN model using both local and global interpretability techniques, specifically, mean activation analysis and correlation-based attribution, which are commonly adopted strategies in explainable deep learning (Stano et al., 2020; Zhang and Zhu, 2018). Considering a test dataset $\boldsymbol{D}_{test} = \left[\mathbf{x}_i \quad q_i^{CHF}\right]_{i=1,\ldots,N^t}$, we calculate the mean activation of each feature map (convolutional or attention) across the test set to quantify their average contributions. The analysis of the activations using the test set reflects how the trained model generalizes.

For a given intermediate layer l in a feature extraction block b, let the output activation tensor be $\mathbf{C}^l \in \mathbb{R}^{N^t \times L_b^l \times f^b}$, where L_b^l is the sequence (feature map) length. Then, the filter-wise mean activation for a filter $f \in \left\{1, 2, \ldots, f^b\right\}$ is:

$$\mu_f^l = \frac{1}{N^t \cdot L_b^l} \sum_{i=1}^{N^t} \sum_{p=1}^{L_b^l} \left| c_{i,p,f}^l \right| \tag{14}$$

Similarly, the position-wise mean activation for feature map spatial position $p \in \{1, 2, \ldots, L_b^l\}$, which capture the information related to the specific physical parameter or their combinations for the i-th prediction, is:

$$\mu_{i,p}^l = \frac{1}{f^b} \sum_{f=1}^{f^b} \left| c_{i,p,f}^l \right| \tag{15}$$

Then, the average mean activation of the feature map position p over the entire test set is:

$$\mu_p^l = \frac{1}{N^t} \sum_{i=1}^{N^t} \mu_{i,p}^l \tag{16}$$

The sample-wise average mean activation for the i-th prediction in a test set is:

$$\mu_i^l = \frac{1}{L_b^l} \sum_{p=1}^{L_b^l} \mu_{i,p}^l \tag{17}$$

Thus, the layer-wise average global mean activation score of layer l in a feature extraction block b is:

$$\mu^l = \frac{1}{N^t} \sum_{i=1}^{N^t} \mu_i^l \tag{18}$$

To assess alignment with physical input parameters and CHF, we compute Pearson correlations between mean activations and the input/output variables. This allows identification of features most influential in model prediction, thereby improving interpretability and scientific trust. Specifically, for an input parameter $\mathbf{x}_j$, the correlation with feature mean activation vector $\boldsymbol{\mu}_p^l$ at position p in layer l of length N^t is:

$$\rho_{\mathbf{x}_j \boldsymbol{\mu}_p^l} = \frac{\mathrm{Cov}\left(\mathbf{x}_j, \boldsymbol{\mu}_p^l\right)}{\sigma_{\mathbf{x}_j} \sigma_{\boldsymbol{\mu}_p^l}} \tag{19}$$

Similarly, the correlation with the aggregated feature mean activation vector $\boldsymbol{\mu}^l$ of layer l within a feature extraction block b is:

$$\rho_{\mathbf{x}_j \boldsymbol{\mu}^l} = \frac{\mathrm{Cov}\left(\mathbf{x}_j, \boldsymbol{\mu}^l\right)}{\sigma_{\mathbf{x}_j} \sigma_{\boldsymbol{\mu}^l}} \tag{20}$$

These correlation metrics facilitate domain-informed global interpretation by quantitatively linking latent feature activations to physical parameters: pressure, tube diameter, mass flux, heated length and outlet quality. This analysis offers critical insights into which input variables predominantly influence the internal feature representations and how this influence evolves across successive layers of the network.

4. Case study

The performance of the proposed SA-1DCNN method is validated

Table 1
Summary statistics of the input parameters of CHF database.

Statistic	Tube Diameter [m]	Heated Length [m]	Pressure [kPa]	Mass Flux [kg/m^{2}s]	Outlet Quality [-]	CHF [kW/m^2]
Min	0.0020	0.050	100.00	8.20	−0.4970	50.00
Max	0.0160	20.000	20000.00	7964.00	0.9990	16339.30
Mean	0.0089	2.625	9681.14	2009.78	0.3558	1808.04
Std. Dev.	0.0025	2.044	5604.56	1511.23	0.2727	1591.65
Q1 (25 %)	0.0079	1.000	5171.00	985.00	0.1580	750.00
Median (50 %)	0.0081	2.000	9800.00	1597.00	0.3290	1424.00
Q3 (75 %)	0.0100	3.658	14000.00	2772.00	0.5343	2333.00

Table 2
Hyperparameters search ranges.

Hyperparameter	Search Range
Number of feature extraction blocks (B)	[1, 4]
Number of filters per block $f^b (b = 1, 2, \cdots, B)$	[16, 512]
Convolutional kernel (filter) size (k)	[2, 5]
Number of nodes in the dense layer (u)	[64, 1024]
Learning rate (η)	[0.0001, 0.01]
Batch size (m)	[32, 512]
Scaling parameter (β)	[1, 6]
Offset (α)	[0, 2]

using the public experimental CHF dataset published by the US NRC along with Ref. (Groeneveld, 2019). The dataset contains data from 59 experiments conducted in vertical water-cooled uniformly heated circular channels in the past decades. This data has been recently made available by the WPRS-EGMUP task force on AI and ML for Scientific Computing in Nuclear Engineering project, a project supported by the OECD/NEA (Le Corre et al., 2024). It contains 24579 data points of five input parameters and one output (CHF), covering a wide range of system operating conditions. The input variables are pressure, tube diameter, mass flux, heated length and outlet quality. Table 1 presents the statistical distributions of the input parameters and CHF values used in this work. Tube diameter ranges from 0.002 to 0.016 m, with most cases concentrated near the median of 0.0081 m. Heated length spans a broad interval (0.05–20 m), though the interquartile range (1.0–3.7 m) indicates clustering in short-to-intermediate lengths, with limited representation of extended channels. System pressure covers 0.1–20 MPa, with a median of 9.8 MPa and balanced coverage across subcooled to near-supercritical conditions; however, cases above 15 MPa are less frequent. Mass flux ranges from 8.2 to 7964 kg/m^{2}s, though most data fall between 985 and 2772 kg/m^{2}s, consistent with typical reactor operating conditions. Outlet quality spans −0.497 to 0.999, which captures subcooled to highly vaporized states, but high-quality (>0.8) cases are comparatively sparse. CHF values range widely (50–16,339 kW/m^2) with a skew toward low-to-intermediate fluxes (median = 1424 kW/m^2). Overall, the database ensures broad parameter coverage, though imbalances remain at distribution extremes, particularly at high mass flux, long heated lengths, and near-dryout conditions.

The data is randomly divided into 80 % (64 % training dataset and 16 % validation dataset) for developing the proposed SA-1DCNN model and 20 % (test dataset) for evaluating the performance of the model, with a fixed random seed to ensure reproducibility. No stratification criteria have been applied; instead, a purely random partitioning strategy has been adopted. Given the broad distribution of operating parameters, this approach is considered sufficient to maintain representative coverage of the experimental conditions across all subsets.

To measure the performance of the SA-1DCNN model on the test dataset $\boldsymbol{D}_{test} = \left[\mathbf{x}_i \quad q_i^{CHF}\right]_{i=1,\ldots,N^t}$, the metrics used in (Ahmed et al., 2025b) are employed, which are:

1) root mean square percentage error (RMSPE), computed as:

$$RMSPE = 100\sqrt{\frac{1}{N^t}\sum_{i=1}^{N^t}\left(\frac{\hat{q}_i^{CHF} - q_i^{CHF}}{q_i^{CHF}}\right)^2} \tag{21}$$

where N^t is the number of observations in $\boldsymbol{D}_{test}$;

2) mean absolute percentage error (MAPE), calculated as:

$$MAPE = \frac{100}{N^t}\sum_{i=1}^{N^t}\left|\frac{\hat{q}_i^{CHF} - q_i^{CHF}}{q_i^{CHF}}\right| \tag{22}$$

3) Q^2-error, computed as:

$$Q^2 = \frac{\sum_{i=1}^{N^t}\left(\hat{q}_i^{CHF} - q_i^{CHF}\right)^2}{\sum_{i=1}^{N^t}\left(\hat{q}_i^{CHF} - \hat{\mu}\right)^2} \tag{23}$$

where $\hat{\mu}$ is the average of the experimental CHF values, $\boldsymbol{q}^{CHF} \in \mathbb{R}^{N^t}$ in $\boldsymbol{D}_{test}$;

4) Average of the ratio of Predicted to Measured CHF values, calculated as:

$$mean\ P \Big/ M = \frac{1}{N^t}\sum_{i=1}^{N^t}\left(\frac{\hat{q}_i^{CHF}}{q_i^{CHF}}\right) \tag{24}$$

5) standard deviation of the ratio of Predicted to Measured CHF values, calculated as:

$$std\ P \Big/ M = \sqrt{\frac{\sum_{i=1}^{N^t}\left(\frac{\hat{q}_i^{CHF}}{q_i^{CHF}} - mean\ P \Big/ M\right)^2}{N^t}} \tag{25}$$

5. Results and discussion

5.1. Analysis and discussion of the CHF prediction results

The optimal architecture of the proposed SA-1DCNN model is set considering its performance on the validation dataset by following the procedures described in Section 3.1. During optimization of the model hyperparameters by BO, the number of epochs is set to 600, and the early stopping is applied with a patience of 30 iterations to prevent overfitting. To further prevent overfitting, a dropout of 0.01 is used for the first and third feature extraction blocks. To train the model, we used the Adam optimization technique (Kingma and Ba, 2015). The pool size for the max pooling operations is set to 2. Optimal hyperparameters are obtained by BO algorithm, with MAPE used as the objective function (Eq. (13)). Table 2 reports the hyperparameters along with their respective search ranges used in the BO framework.

The hyperparameter vector, $\boldsymbol{\theta} = \left(B, f^1, f^2, f^3, k, u, \eta, m, \beta, \alpha\right)$ explored by the BO search and the corresponding objective function, $\mathscr{L}_{BO}(\boldsymbol{\theta})$, are shown in Fig. 3 through Fig. 6 over 100 iterations. In each iteration, the

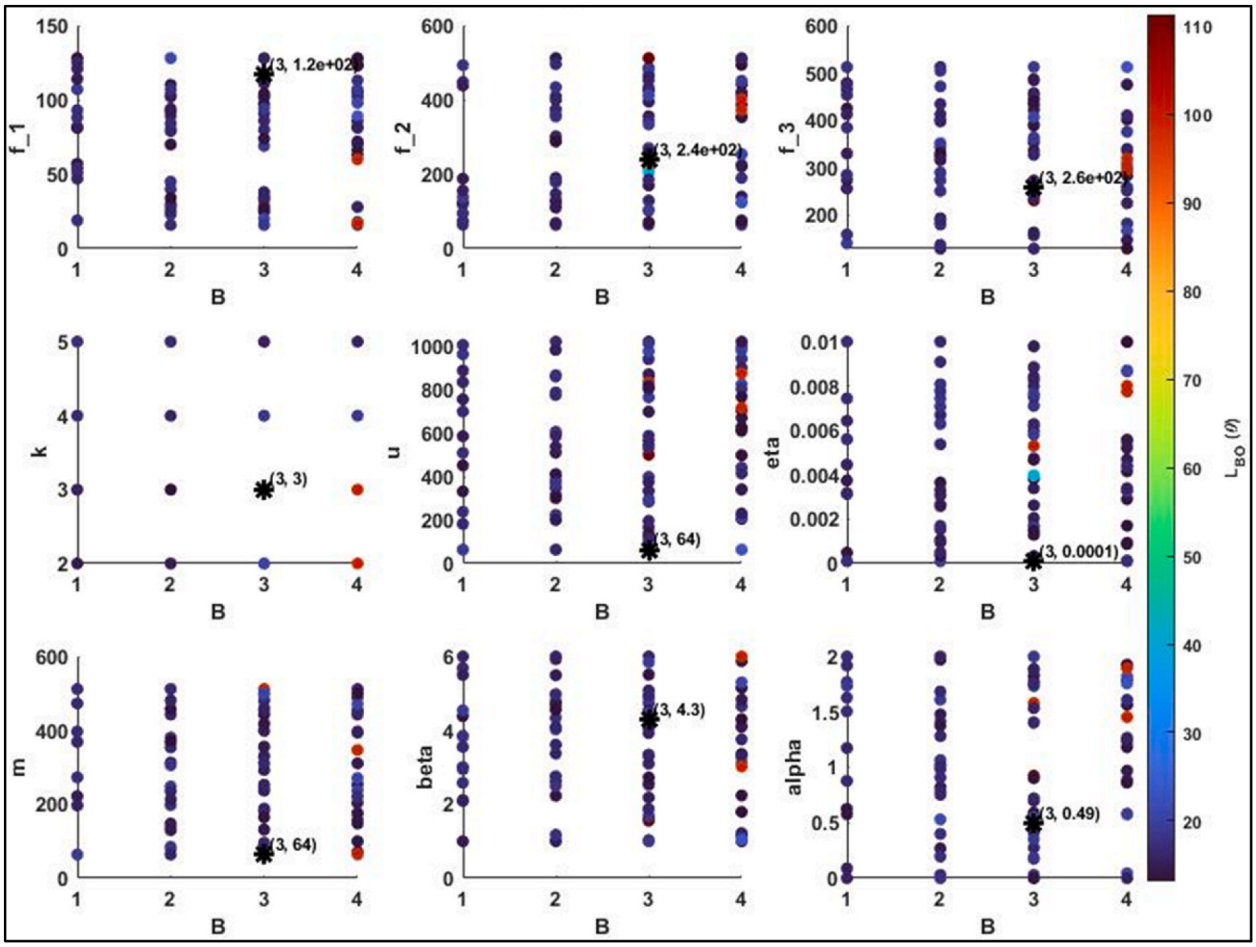

Fig. 3. Hyperparameter values explored by the BO search: plots of the hyperparameters against B in the solutions and associated $L_{BO}(\theta)$.

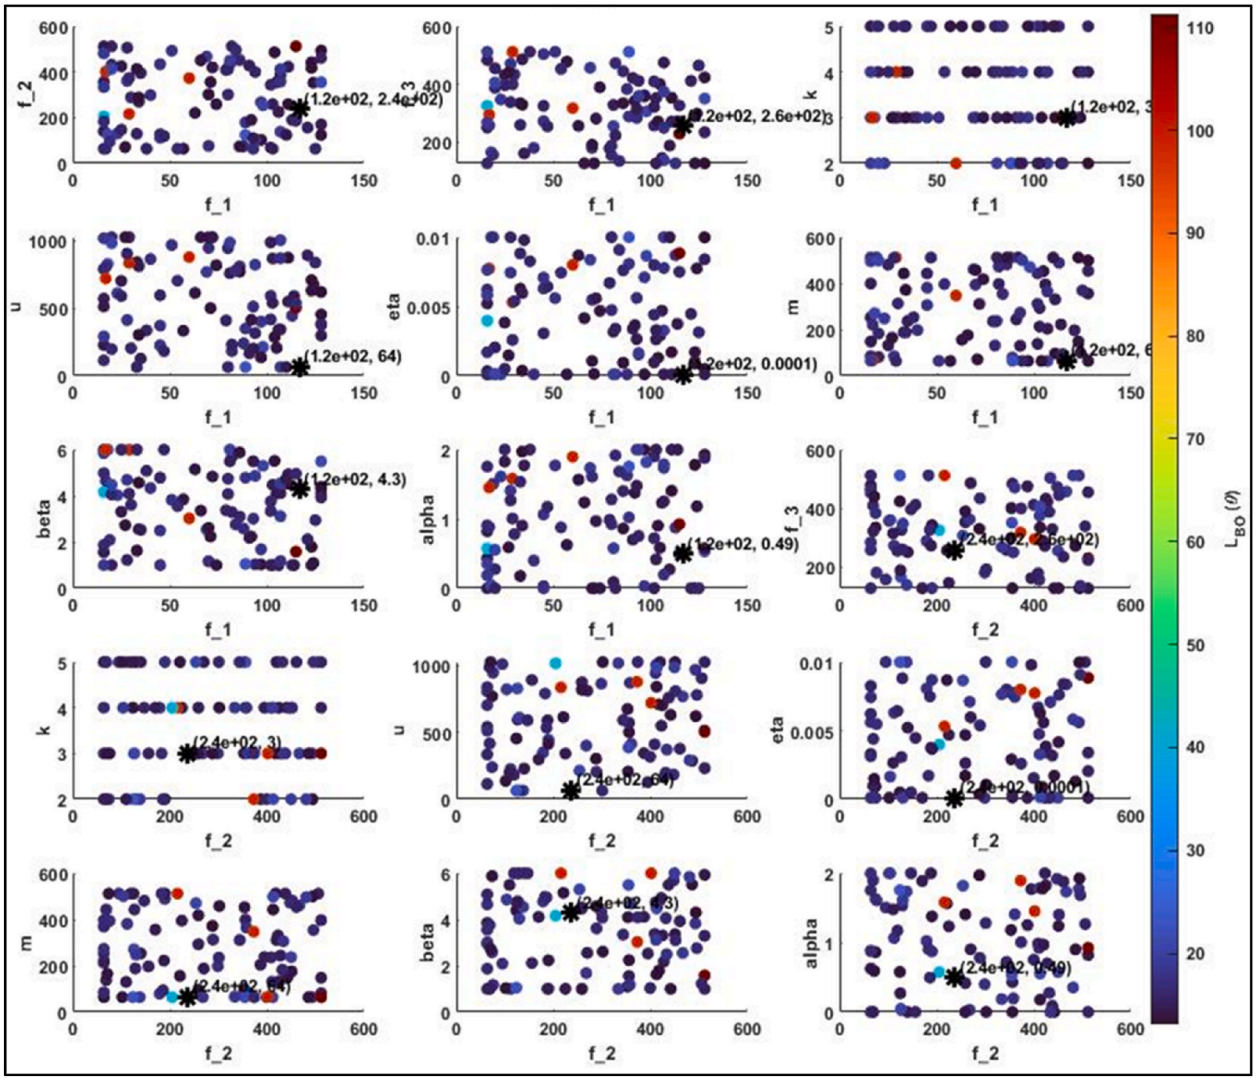

Fig. 4. Hyperparameter values explored by the BO search: plots of the hyperparameters against f^1 and f^2 in the solutions and associated $L_{BO}(\theta)$.

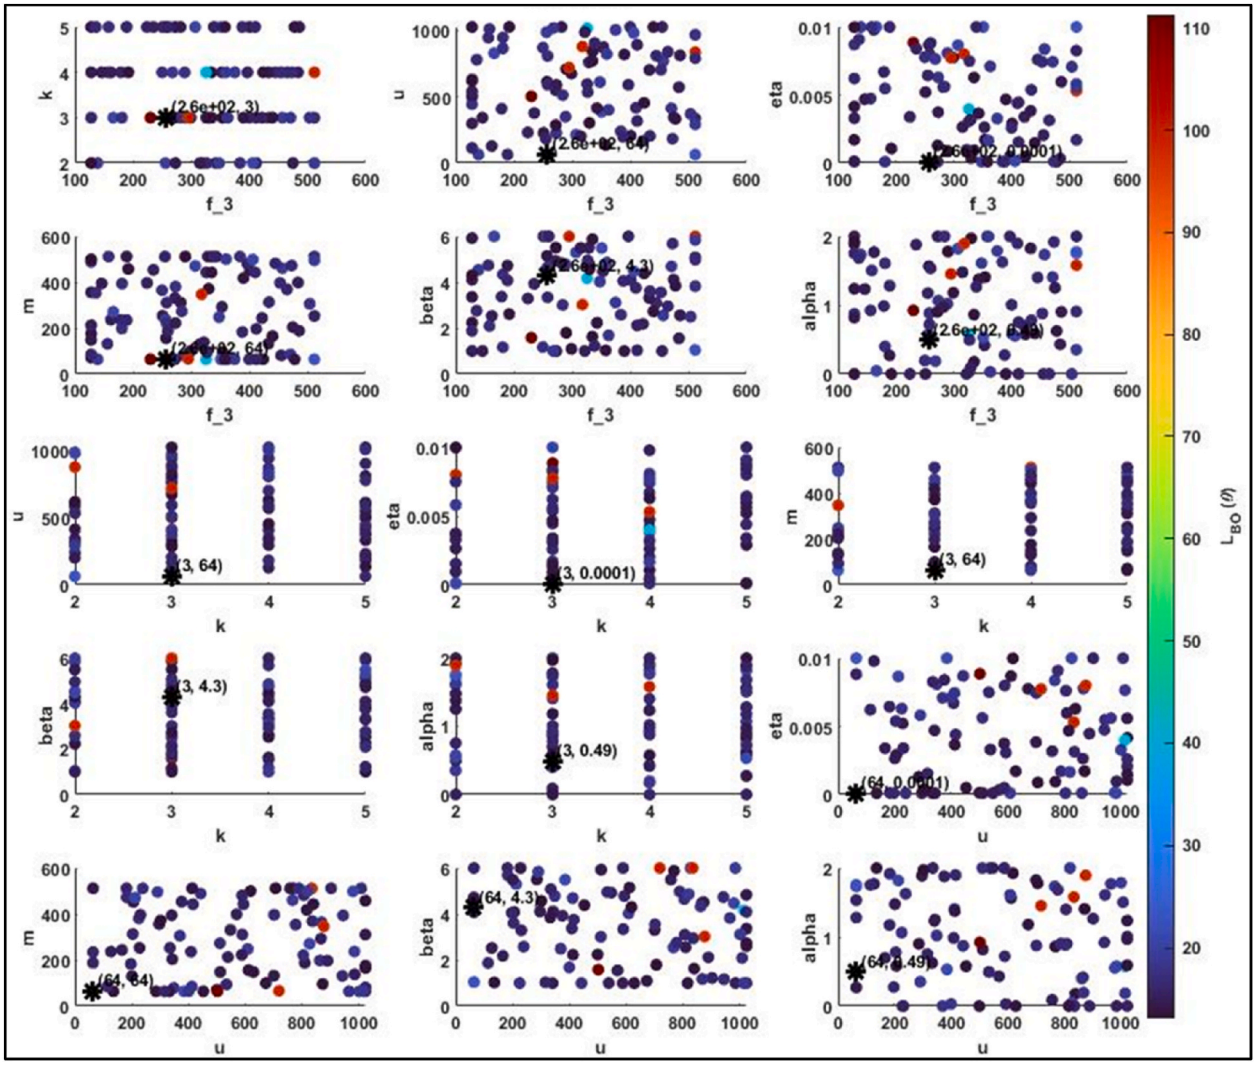

Fig. 5. Hyperparameter values explored by the BO search: plots of the hyperparameters against f^3, k and u in the solutions and associated $L_{BO}(\theta)$.

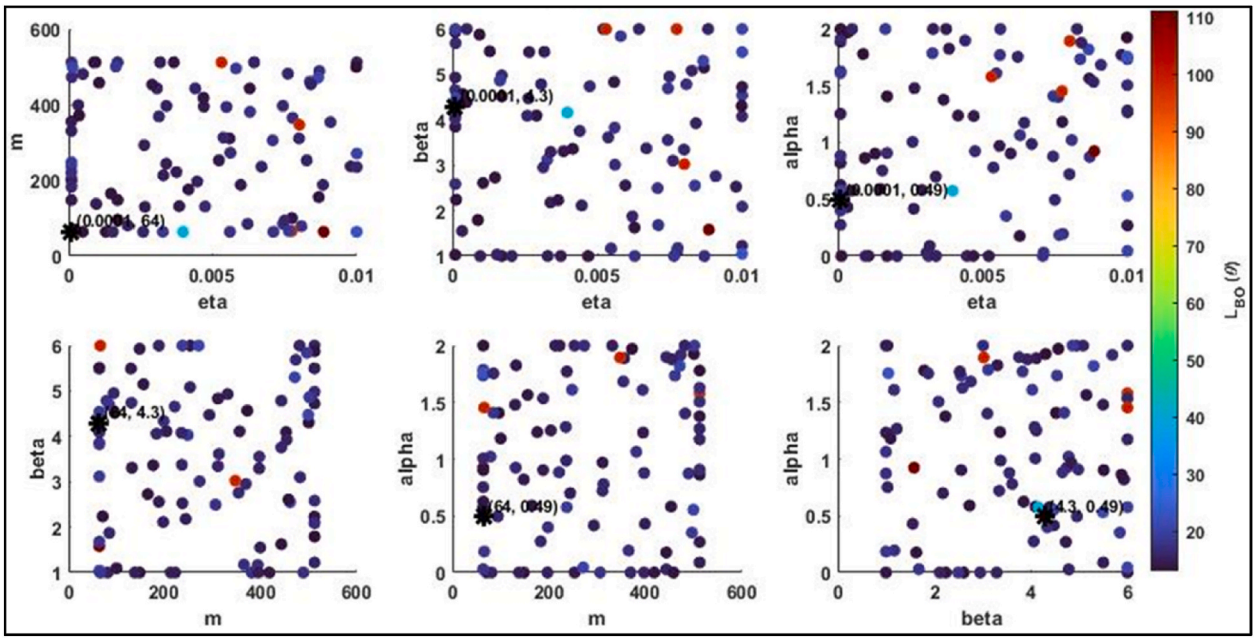

Fig. 6. Hyperparameter values explored by the BO search: plots of the hyperparameters against η, m, β and α in the solutions and associated $L_{BO}(\theta)$.

BO algorithm provides values for the hyperparameters $\boldsymbol{\theta}$ of the SA-1DCNN model, which are, then, used to configure and train a SA-1DCNN model using the training dataset. The performance of the trained model is subsequently evaluated on the validation set to compute $\mathscr{L}_{BO}(\boldsymbol{\theta})$. Fig. 3 presents scatter plots of the hyperparameters $\left(f^1, f^2, f^3, k, u, \eta, m, \beta, \alpha\right)$ against the hyperparameter B (i.e., the number of feature extraction blocks) alongside the associated $\mathscr{L}_{BO}(\boldsymbol{\theta})$ values. It is evident from Fig. 3 that when B exceeds 3, the optimization function ($\mathscr{L}_{BO}(\boldsymbol{\theta})$) generally does not exhibit significant improvement with variations in the values of other hyperparameters. This suggests that increasing the number of feature extraction blocks beyond three does not necessarily enhance performance and may, in fact, significantly increase computational burden and potentially lead to overfitting. Notice that when B is within the interval $[1, 3]$ (Fig. 3), the optimization function ($\mathscr{L}_{BO}(\boldsymbol{\theta})$) remains relatively stable despite changes in the values of convolutional filters (f^1, f^2, f^3). This indicates minimal sensitivity of the model performance to the number of filters in this region. Further insights from Figs. 4–6 reveal that fixing the value of any of the convolutional filters (f^1, f^2, f^3) while increasing the others generally leads to improved optimization performance. Additionally, it is noteworthy that when a very small value of η (learning rate) is fixed, the optimization performance $\mathscr{L}_{BO}(\boldsymbol{\theta})$ tends to improve with variations in other hyperparameters, suggesting that fixing a relatively small value for η can support effective optimization. Interestingly, the filter/kernel size k appears to have a random influence on performance when the number of filters (f^1, f^2, f^3) is varied, thus highlighting its importance and sensitivity in achieving optimal model configuration. Moreover, it is observed that variations in the fully-connected layer size u, batch size m, and the output activation parameters β (scaling) and α (offset) tend to exert a non-systematic, random influence on $\mathscr{L}_{BO}(\boldsymbol{\theta})$ in interaction with other hyperparameters. This randomness underscores the relevance of these parameters in model optimization, emphasizing the need for careful tuning to ensure optimal performance of the proposed model. Table 3 reports the obtained set of optimal hyperparameters, $\boldsymbol{\theta}$ at the minimum value of $\mathscr{L}_{BO}(\boldsymbol{\theta}) = 13.1323$. These values are respectively displayed on each of the plots in Fig. 3 through Fig. 6.

The performance of the trained SA-1DCNN model is evaluated on the test dataset through a comparative analysis with four state-of-the-art CHF prediction approaches: a traditional LUT, Support Vector Regression (SVR), Random Forest (RF) and an optimized ensemble of NNs proposed in (Ahmed et al., 2025b). Both the SVR and RF models are developed and optimized using the same 80 % training dataset employed for the development of the proposed SA-1DCNN method. The optimized ensemble of NN model is trained using the same dataset and by following the same configuration described in (Ahmed et al., 2025b). Fig. 7 presents a comparison between the CHF predictions of the proposed and LUT methods on the test dataset. As shown in Fig. 7, a larger amount of CHF datapoints predicted by the proposed model (92.9 %) fall within a ±20 % error band than those from the LUT (68.7 %), which shows that the proposed SA-1DCNN method is effective in capturing the relationships between input parameters and CHF with robust generalization capabilities. Table 4 reports the values of the performance metrics obtained by each method. The results show that the proposed SA-1DCNN model significantly surpasses all other methods, which demonstrates superior performance across all metrics. Specifically, the

Table 3
Optimal hyperparameters obtained for the proposed model using BO algorithm.

Hyperparameter	Optimal value
Number of feature extraction blocks (B)	3
Number of filters per block $f^b (b = 1, 2, 3)$	117, 238, 257
Convolutional kernel (filter) size (k)	3
Number of nodes in the dense layer (u)	64
Learning rate (η)	0.0001
Batch size (m)	64
Scaling parameter (β)	4.2916
Offset (α)	0.4913

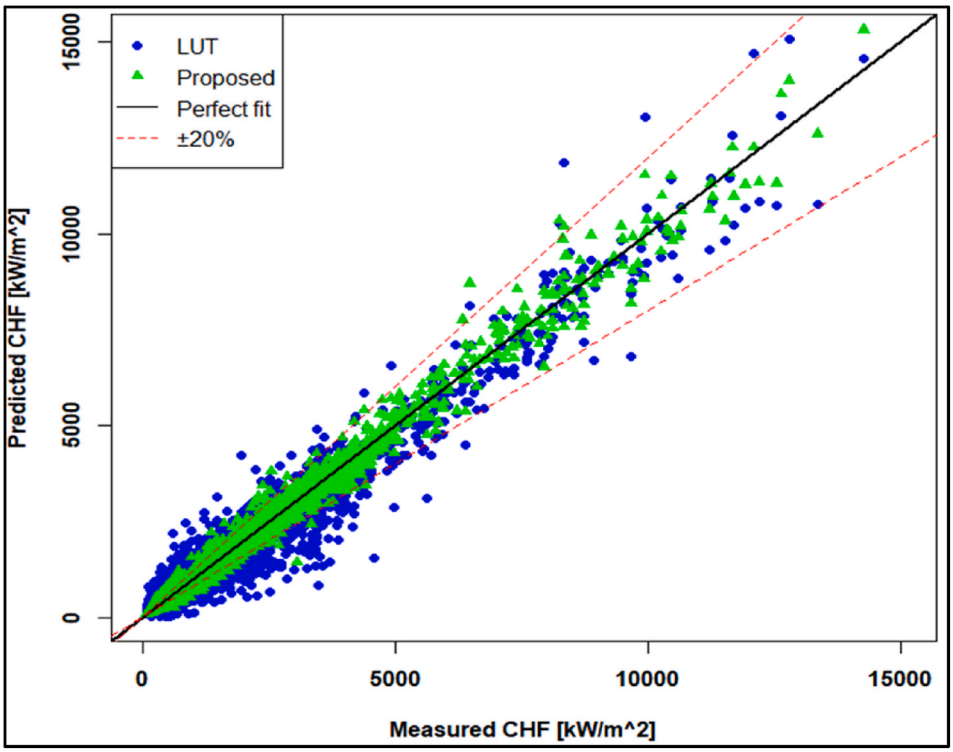

Fig. 7. Plot of predicted CHF against measured CHF on test dataset for the proposed and LUT methods.

proposed SA-1DCNN model achieves the lowest MAPE and RMSPE, alongside a near-zero Q^2-error, which signify a strong predictive capability. The SA-1DCNN model significantly improves prediction accuracy, with 92.9 % of the predicted datapoints falling within a 20 % error band, a performance superior to that of the other models. Furthermore, the average prediction deviation is only +0.4 % relative to the actual values, which is significantly lower than +1.1 % observed for the optimized ensemble of NN model (the best-performing state-of-the-art method). Also, the proposed SA-1DCNN model achieves the lowest standard deviation of the P/M ratio (0.114), which reflects enhanced consistency and robustness in the predictions of the proposed method.

To further validate the proposed SA-1DCNN model, the predictions of the entire CHF dataset (24579 datapoints) are used to compute its performance metrics and compare against state-of-the-art AI-based models: the transformer and CNN models developed in (Zhou et al., 2024), the benchmark NN model in (Le Corre et al., 2024) and the most recently published optimized ensemble of NN model proposed in (Ahmed et al., 2025b). This facilitates a fair and consistent comparison between the proposed approach and the AI-based models presented in (Zhou et al., 2024; Ahmed et al., 2025b), as all models were evaluated using the same 24579 CHF datapoints. The comparative performance metrics across the entire CHF dataset are summarized in Table 5. The proposed model consistently outperforms other AI-based methods across all evaluation metrics, which achieves a RMSPE of 9.44 %, a MAPE of 5.77 %, a Q^2-error of 0.99 %, a mean P/M of 1.002 and a std P/M of 0.094.

The proposed method achieves substantial improvements in CHF prediction accuracy over existing approaches. Specifically, it reduces the MAPE by approximately 15 % and 71 % when compared to the best-performing (optimized ensemble of NNs) and worst-performing (LUT) comparison methods, respectively. Similarly, in terms of RMSPE, the proposed model shows a reduction of approximately 10 % and 71 % relative to the best-performing (optimized ensemble of NNs) and worst-performing (LUT) comparison methods, respectively. With respect to

Table 4
Performance results of the CHF prediction models on the test dataset.

Model	Performance on Test Dataset					
	RMSPE	MAPE	Q^2	Data within ±20 %	Mean P/M	Std. P/M
LUT	36.19	20.52	6.08 %	68.71 %	1.042	0.359
RF	13.47	8.79	1.98 %	90.23 %	1.013	0.134
SVR	20.75	12.28	3.79 %	85.31 %	1.024	0.206
Optimized Ensemble of NNs (Ahmed et al., 2025b)	12.42	7.92	1.50 %	91.81 %	1.011	0.122
Proposed SA-1DCNN Model	**11.41**	**7.23**	**1.32 %**	**92.90 %**	**1.004**	**0.114**

Table 5
Performance results of the CHF prediction models on the entire CHF dataset.

Model	Performance on the entire 24579 datapoints					
	RMSPE	MAPE	Q^2	Data within ±20 %	Mean P/M	Std. P/M
LUT	36.30	19.77	6.31 %	68.90 %	1.032	0.362
Benchmark NN (Le Corre et al., 2024)	12.60	8.03	2.20 %	–	1.010	0.115
NN (Zhou et al., 2024)	20.40	11.20	1.98 %	–	1.034	0.201
CNN (Zhou et al., 2024)	16.80	9.72	1.36 %	–	1.027	0.166
Transformer (Zhou et al., 2024)	12.30	7.22	1.26 %	–	1.008	0.122
Optimized Ensemble of NNs (Ahmed et al., 2025b)	10.43	6.76	1.46 %	94.38 %	1.007	0.104
Proposed SA-1DCNN Model	**9.44**	**5.77**	**0.99 %**	**95.60 %**	**1.002**	**0.094**

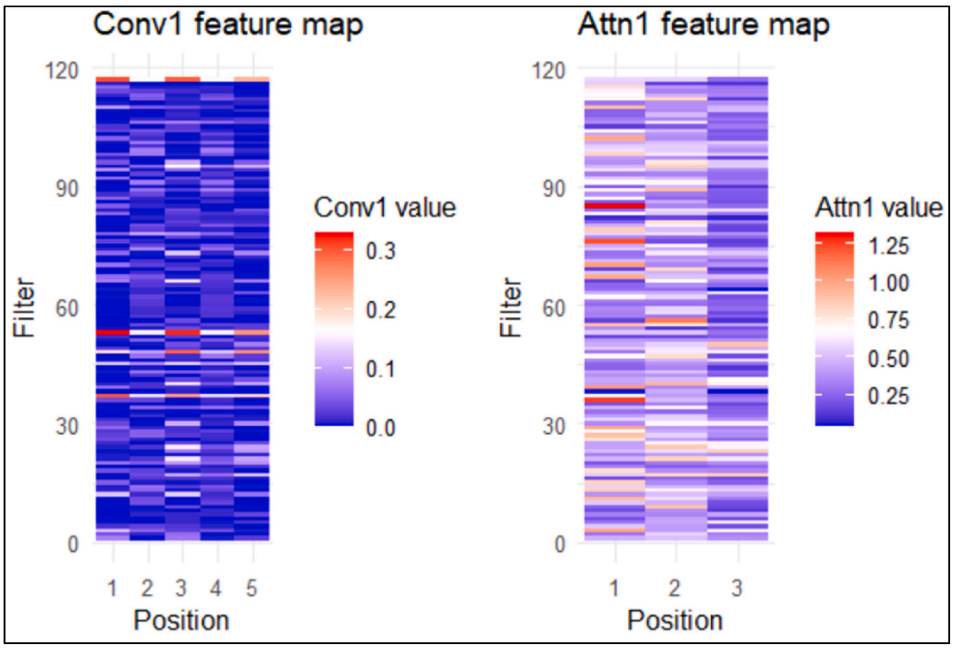

Fig. 8. Visualization of the mean activations of convolutional and self-attention layers features of the first block based on the test set.

Q^2-error, the model achieves a reduction of approximately 21 % and 84 % when compared to the best-performing (Transformer) and worse-performing (LUT) comparison methods, respectively. These results clearly demonstrate the superior predictive accuracy of the proposed SA-1DCNN model, thereby establishing it as the most accurate CHF prediction method for overall prediction performance to date. The enhanced performance is attributed to the integration of a self-attention mechanism within the 1DCNN. This mechanism allows the model to dynamically assess the importance of input features across the parameter space, thereby improving its ability to capture complex, nonlinear relationships. Such relationships, which often involve long-range dependencies and non-local interactions, are difficult to learn using conventional models. The proposed attention-enhanced 1DCNN benefits from both local pattern extraction and global context awareness, resulting in more robust and generalizable predictions.

Besides the overall superior performance, the proposed model exhibits a maximum MAPE of 120 % for individual datapoints in the CHF database, a notable improvement over the 160 % maximum MAPE reported in (Zhou et al., 2024) for the Transformer model. Among the 24,579 datapoints, only 413 datapoints have a MAPE exceeding 30 %, which are significantly fewer than the 801 datapoints reported in (Zhou et al., 2024) for the Transformer model. Even more, the LUT model records a maximum MAPE of 2209 %, with 4541 datapoints having a MAPE greater than 30 %. These results further underscore the reliability and robustness of the proposed model across a wide range of test conditions.

5.2. *Explainability and interpretability analyses*

To explain the behavior of the proposed model on unseen data for reflecting real-world performance, the data of the test set is used to prevent any biased interpretation. To gain insights into the feature extraction behavior of the proposed SA-1DCNN model, we first examine the mean activations of the first convolutional layer (Conv1) and the first self-attention layer (Attn1). The mean activations are computed across all test set data for each spatial position. These positions correspond to the five physical input parameters: pressure, tube diameter, mass flux, heated length and outlet quality. Fig. 8 shows the distributions of the mean activations for the features of these two layers in the first feature extraction block (Block #1). Note that, based on the architecture of the SA-1DCNN model (Fig. 2), which incorporates zero-padding and the obtained optimal hyperparameters in Table 3, the output of Conv1 is $\mathbf{Z}^1 \in \mathbb{R}^{N^t \times 5 \times 117}$, with five spatial positions (each corresponding to one of the five input parameters) and 117 convolutional filters, which are the learned local feature maps that capture interdependences among the physical input parameters. Also note that the self-attention mechanism is applied after batch normalization and max pooling (Fig. 2), resulting in an output $\mathbf{F}^1 \in \mathbb{R}^{N^t \times 3 \times 117}$, with three positions each corresponding to the combination of two positions (pool size $= 2$) from the five positions of Conv1 layer, which reduces the spatial resolution from five input positions to three pooled segments. It can be observed from Fig. 8 that the Conv1 activations show slightly higher sensitivity to the third position, which corresponds to pressure. This may indicate that pressure-related patterns contribute more prominently to the early learned representations. The relatively uniform activations across the other positions suggest that the Conv1 filters respond consistently to localized variations in the input parameters, without overly emphasizing any single feature. This behavior aligns with the intended function of early convolutional layers in the deep architecture (i.e., to capture local dependencies and spatially structured features across all inputs without immediate global context).

Instead, the self-attention mechanism dynamically re-weights the intermediate features by emphasizing relevant spatial interactions and contextual relationships among the inputs. The relatively higher attention scores at the first and second positions indicate that the model assigns more relevance to the first two pooled regions, which collectively span the first four input positions (i.e., tube diameter, heated length, pressure and partially mass flux). On the contrary, the last region receives lower attention, indicating a reduced emphasis on outlet quality and possibly part of mass flux in the global contextualization. This suggests that the model learns to focus more on geometric (tube diameter, heated length) and thermodynamic parameters (pressure) in the early attention blocks, likely due to their higher contribution to the accurate prediction of CHF. These findings support a coherent transition from uniform local extraction (Conv1) to selective global integration (Attn1), where the attention mechanism enhances convolutional features by incorporating broader context and non-local dependencies, thus improving the interpretability and predictive capability of the model.

To further investigate model behavior under different physical

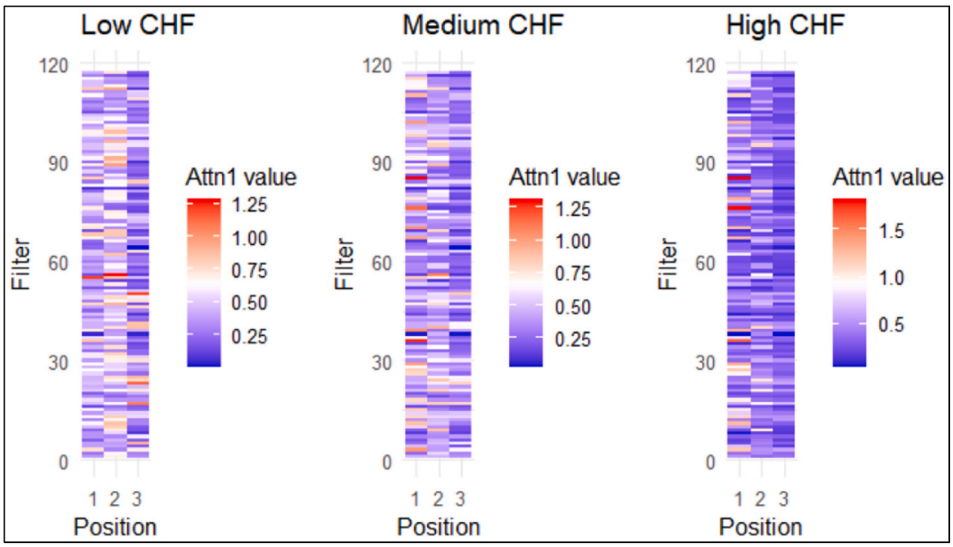

Fig. 9. Visualization of the mean activations of the self-attention layer (attn1) of the first block based on CHF groups of the test set.

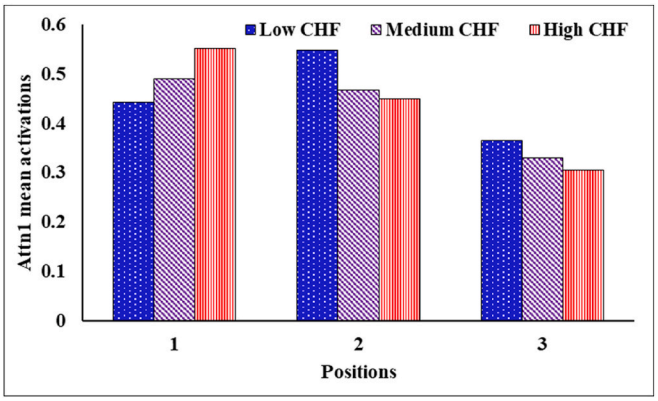

Fig. 10. Summary of the average of the mean activations of the self-attention layer (attn1) of the first block based on CHF groups of the test set.

conditions, we stratified the test set into three CHF intensity categories, low, medium and high, based on the lower and upper tertiles of measured CHF values (i.e., quantiles at 1/3 and 2/3). Fig. 9 shows the visualization of each group for Attn1 layer, which has been summarized in Fig. 10. Notice that the mean activation of the first attention position increases monotonically across the CHF strata, suggesting that geometrical parameters (tube diameter and heated length) become more salient at higher CHF regimes, where channel geometry governs thermal limits (Todreas and Kazimi, 2011). Conversely, the second attention region, dominated by pressure and mass flux, shows the highest activation in the low CHF group, decreasing with higher CHF levels. This trend aligns with two-phase flow physics, wherein pressure and mass flux are crucial at low CHF due to their role in subcooling and flow stability. The third attention region, largely influenced by outlet quality, shows a declining activation trend with rising CHF. This is physically consistent: vapor quality significantly affects low CHF by initiating dryout, but becomes less dominant at high CHF, where flow stability and geometric constraints take precedence. These stratified attention behaviors reveal that the proposed SA-1DCNN model dynamically adjusts its focus across CHF regimes, emphasizing vapor quality and pressure in low CHF scenarios, and geometry in high CHF ones. This not only corroborates domain knowledge but also enhances the interpretability and physical credibility of the model.

To assess whether the proposed explainable SA-1DCNN model learns relationships that are consistent with established CHF mechanistic theories, we analyzed the group-specific mean activations of the first attention layer (Attn1) across low, medium and high ranges of each input parameter in the test set. This allows demonstrating the alignment with established mechanistic theories of CHF. Unlike empirical correlations, which often lack explicit physical justification, mechanistic CHF theories attribute burnout to well-defined thermohydraulic processes. Fig. 11 summarizes the mean group-wise activations of Attn1 across the five input parameters. The attention positions 1, 2 and 3 correspond to geometry (channel diameter and heated length), forcing (pressure and mass flux) and quality (outlet quality), respectively. Thus, the reported mean activations reflect how the model "attends" to geometry (position 1), thermohydraulic forcing (position 2), and exit quality/flow-regime effects (position 3) under different operating conditions.

Several trends observed in Fig. 11 are consistent with mechanistic expectations of CHF behavior. For tube diameter (Fig. 11(a)), attention activations at position 1 increase strongly in the high-diameter group, whereas positions 2 and 3 remain comparatively stable. This trend agrees with the mechanistic understanding that large diameters promote liquid rewetting and delay dryout, thereby reducing CHF sensitivity to flow instabilities (Tong and Tang, 1997). The increased attention weights suggest that the model emphasizes diameter information in regimes where geometry-driven flow distribution is critical. These observations are consistent with hydrodynamic-instability scaling and film-dryout perspectives (Hewitt and Hall-Taylor, 1970; Zuber, 1959). For heated length (Fig. 11(b)), attention values at position 1 decrease

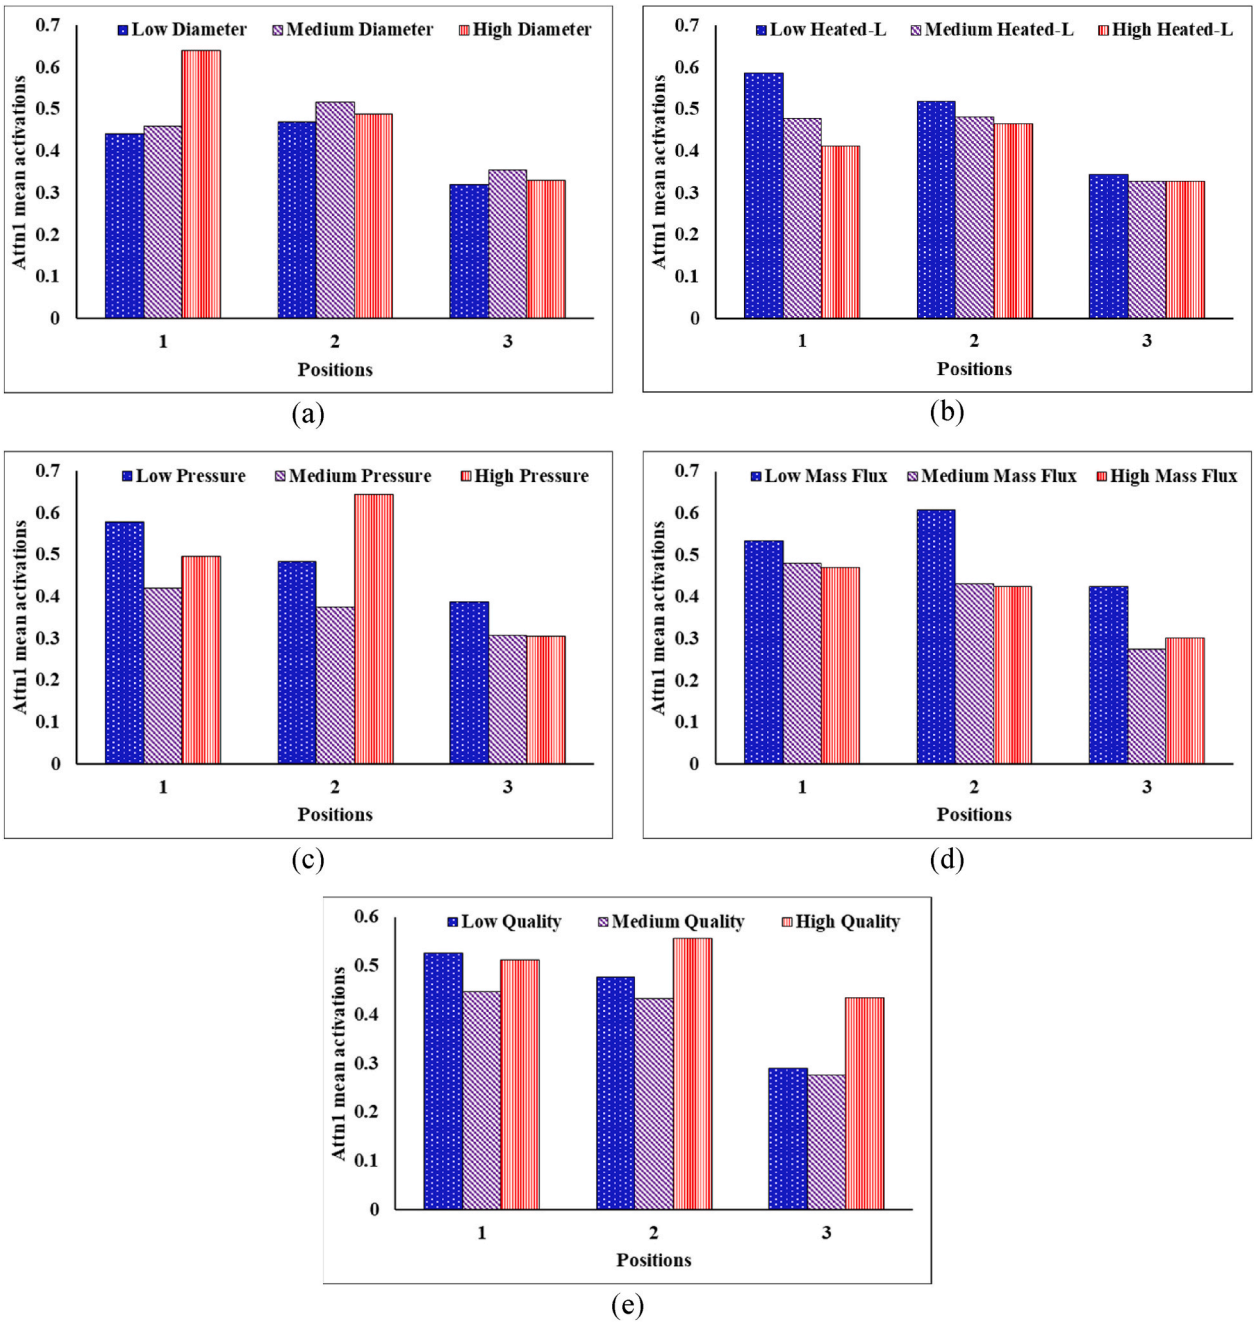

Fig. 11. Summary of the average of the mean activations of the self-attention layer (attn1) of the first block based on: (a) tube diameter groups; (b) heated length groups; (c) pressure groups; (d) mass flux groups; and (e) outlet quality groups of the test set.

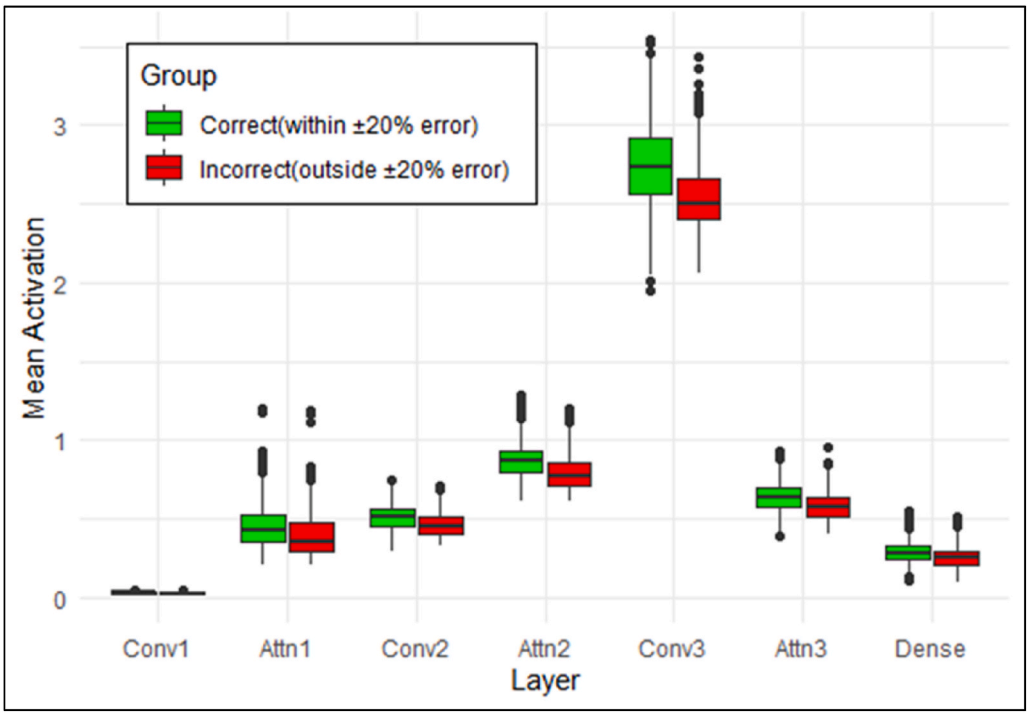

Fig. 12. Distributions of the mean activations across layers of the trained model for correctly and incorrectly predicted CHF groups in the test set.

monotonically from low to high heated lengths, in line with the expectation that CHF declines as film depletion accumulates along longer surfaces, thereby promoting earlier dryout (Tong and Tang, 1997; Weisman and Pei, 1983). The weak variations at positions 2 and 3 are consistent with the notion that CHF decreases with heated length primarily via liquid film exhaustion rather than a fundamental regime change.

For pressure ((Fig. 11(c))), attention distributions are non-monotonic, with high-pressure cases showing strong activation at position 2. This aligns with mechanistic theories predicting that at elevated pressures, increased vapor density reduces buoyancy while amplifying through-flow and pressure-driven influences on bubble detachment and film stability (Ishii and Mishima, 1984; Zuber, 1959). The decline with pressure at position 3 is also reasonable, as quality effects are less pronounced when nucleation dynamics are pressure-stabilized. For mass flux (Fig. 11(d)), position 2 activations are highest at low mass flux, and substantially lower at medium and high values. This matches the mechanistic knowledge that CHF is highly sensitive to mass flux in the low range, where reduced convective inertia promotes early dryout, whereas at high fluxes CHF increases nearly linearly with diminished sensitivity (Hewitt and Hall-Taylor, 1970; Tong and Tang, 1997; Weisman and Pei, 1983). The emphasis of the model reflects this mechanistic asymmetry.

Finally, for outlet quality (Fig. 11(e)), activations increase distinctly at position 2 in the high-quality group, with a concurrent rise at position 3. This is consistent with the mechanistic expectation that CHF is strongly governed by quality in high ranges, where thinning of the annular liquid film and dryout initiation dominate (Hewitt and Hall-Taylor, 1970; Ishii and Mishima, 1989; Morse et al., 2021). The co-elevation of positions 2 and 3 encodes the coupled mass-flux–quality physics: higher mass flux can delay dryout via deposition but also promote entrainment at very high values. The attention of the model reflects this nonlinear mechanistic interaction.

Although attention does not imply causation, the observed regime-aware shifts across groupings support mechanistic consistency. The analysis validates several CHF characteristics: *a*) mechanistic partitioning emerges naturally, with the model associating position 1 with geometry, position 2 with forcing and position 3 with quality, in agreement with physical expectations; *b*) monotonic and regime-aware signals (e.g., rising position 1 with diameter, falling position 1 with heated length, dominance of position 2 at high pressure and low mass flux, strong position 3 at high quality) qualitatively align with hydrodynamic instability, dry-spot/macrolayer and annular dryout/entrainment frameworks; and *c*) coupled interactions are captured, such as the co-increase of positions 2 and 3 at high quality, which reflects mass-flux–quality interplay central to entrainment–deposition models. Overall, the group-specific attention analysis demonstrates that the proposed SA-1DCNN model not only achieves predictive accuracy but also learns parameter sensitivities and interactions consistent with CHF mechanistic theories. This alignment strengthens the physical credibility of the model, moving it beyond a black-box regression toward an interpretable surrogate capable of supporting safety-critical reactor applications, including licensing, power uprates and higher burn-up operation, where mechanistic justification is essential for regulatory acceptance.

To explore how internal activations relate to predictive accuracy, we compute mean layer activations for each test sample and group them based on prediction correctness. A prediction is here considered correct if the P/M CHF ratio falls within ±20 % of unity. Activation values are extracted using intermediate models and averaged over spatial and filter dimensions for convolutional and attention layers, and directly for the dense layer. The results, summarized in Fig. 12, show clear activation separability between correct and incorrect predictions, particularly in deeper layers. This indicates that stronger and more structured activations are associated with correct predictions, supporting the notion that the internal representations of the model become increasingly discriminative as accuracy improves. These insights could potentially be applied in downstream decision-making tasks, such as model confidence-based acceptance or rejection, although such applications are beyond the scope of this study.

Further interpretability is gained by examining Pearson correlations between mean activations at Conv1 and Attn1 with both the input parameters and measured/predicted CHF values. As shown in Fig. 13, Conv1 Position 1 displays a strong positive correlation with tube diameter and moderate correlations with measured and predicted CHF,

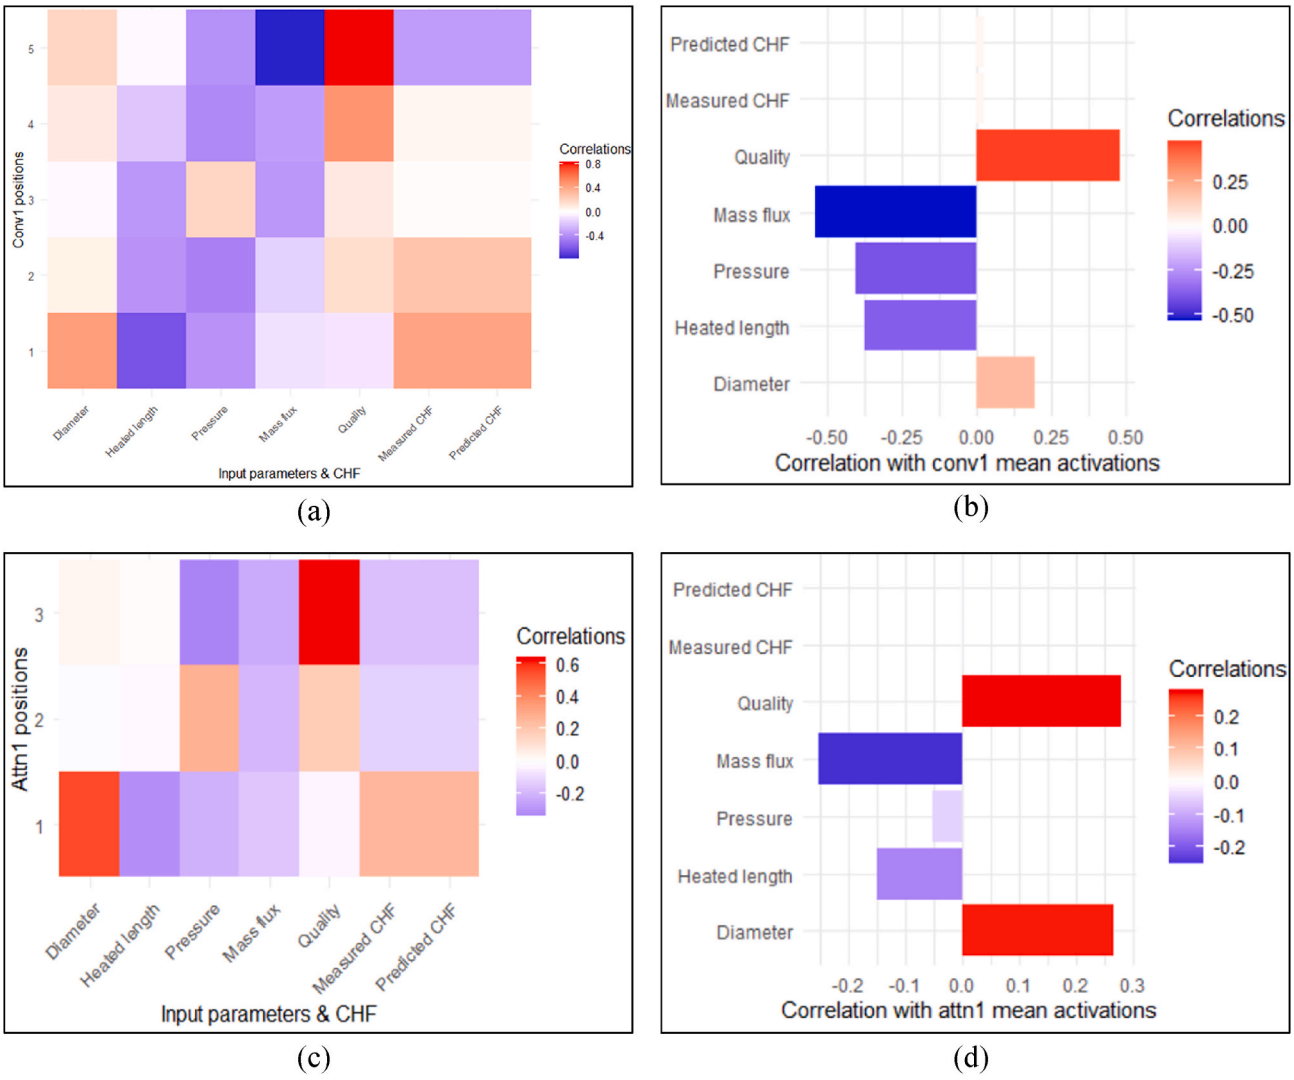

Fig. 13. Correlations between mean activations of each positions of Conv and Attn1 layers of block 1 with input parameters, along with the measured and predicted CHF.

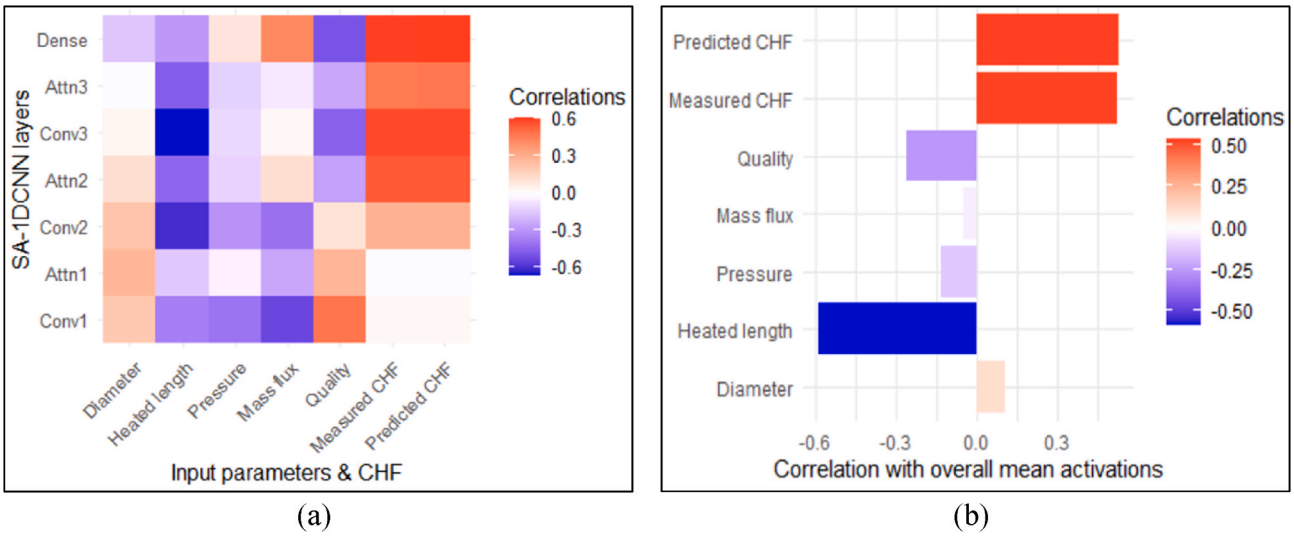

Fig. 14. Correlations between the overall mean activations of each SA-1DCNN layer and the input parameters, along with the measured and predicted CHF.

Table 6
Ranges of slice datasets (Le Corre et al., 2024).

Slice #	Tube Diameter [m]	Heated Length [m]	Pressure [kPa]	Mass Flux [kg/m^2/s]	Outlet Quality [-]
1	**0–0.016**	6.000	14701	998.5	0.391
2	**0–0.016**	6.000	9807	1003.3	0.529
3	0.00801	**0–20**	9806	1000.0	0.587
4	0.00811	**0–20**	2009	752.2	0.756
5	0.00800	0.998	**0–20000**	2040.2	0.140
6	0.01340	3.658	**0–20000**	2040.2	0.378
7	0.00800	1.570	12750	**0–8000**	0.144
8	0.01000	4.966	16000	**0–8000**	0.343
9	0.00814	1.943	9831	1519.5	**−0.5–1.0**
10	0.00800	0.997	17650	2002.7	**−0.5–1.0**

thus demonstrating that early filters are capturing geometry-related effects. In contrast, Position 5 shows strong positive correlation with outlet quality and a significant negative correlation with mass flux, consistent with the physical tradeoff where increasing vapor quality reduces CHF, whereas higher mass flux enhances it via improved cooling (Saha and Zuber, 1974). The aggregate Conv1 mean activations (Fig. 13 (b)) show dominant correlation with outlet quality and inverse correlation with pressure and mass flux, thus reaffirming the role of the convolutional layer in capturing localized, physically meaningful relationships. For Attn1 (Fig. 13(c)), the first attention position is strongly correlated with tube diameter and moderately with CHF. This suggests that learned representations focused on geometrical relevance in this layer. The third attention shows a high positive correlation with outlet quality and a high negative correlation with measured/predicted CHF, which supports the inverse CHF-to-vapor quality relationship. The Attn1 mean activation (Fig. 13(d)) correlates most strongly with outlet quality and shows moderate negative correlation with mass flux. This indicates that attention layers adjust their focus in response to inter-parameter dynamics, beyond simple localized encoding.

These results highlight the complementary roles of convolution and attention mechanisms: Conv1 captures localized, parameter-specific interactions, whereas Attn1 captures global, interaction-aware representations. For instance, Conv1 Position 5 and Attn1 Position 3, both showing strong correlations with outlet quality, underline the critical role of vapor content in CHF prediction. Also, the alignment between specific high-correlation positions and the CHF output reinforces the physical relevance of the learned features and strengthens confidence in the generalization capabilities of the proposed SA-1DCNN model. This interpretable, hybrid feature extraction supports the integration of domain knowledge and deep learning that is crucial for safety-critical applications, such as nuclear reactor heat transfer systems.

To further ensure transparency and physical consistency in the proposed SA-1DCNN model, we performed a correlation-based interpretability analysis by evaluating the Pearson correlation coefficients between the overall mean activations of each intermediate layer and the input physical parameters, as well as with both predicted and measured CHF values. The correlation values are visualized in Fig. 14 along with a second-level analysis of the overall mean activation vector (combined across all layers). The results shown in Fig. 14(a) indicate a progressive abstraction from low-level physical input encoding to high-level task-oriented representation, aligning with the hierarchical feature learning property of deep neural architectures. Early convolutional layers (e.g., Conv1) show moderate correlations with pressure, mass flux and outlet quality, but weak association with the CHF output. This implies that Conv1 captures localized thermodynamic gradients and flow patterns, consistent with the literature on two-phase heat transfer phenomena (Tong and Tang, 1997). The Attn1 layer shows slightly decreased correlation magnitudes, but it begins to shift focus toward relational patterns among inputs rather than their individual values. For example, while Attn1 maintains moderate correlation with outlet quality, its sensitivity to pressure and mass flux is reduced. This indicates a context-aware filtering mechanism that facilitates cross-parameter interactions. Notice that in the deeper layers (e.g., Conv2, Attn2, Conv3, Attn3), a transition is observed. Conv2 maintains strong sensitivity to heated length and mass flux, whereas Attn2 starts aligning closely with CHF targets. Notice that Conv3 and Attn3 continue this alignment by showing high correlations with measured/predicted CHF, while reducing dependence on raw inputs. This indicates that mid-to-deep layers, particularly attention ones, effectively model nonlinear dependencies and latent physical regimes relevant to CHF onset. The final dense layer shows the strongest correlation with predicted and measured CHF, which confirms that the learned representations converge toward task-relevant abstractions. Interestingly, it also reverses the sign of correlations with outlet quality, thus suggesting complex, possibly non-monotonic interactions in the learned features, an expected behavior in boiling crisis conditions where higher outlet qualities lead to lower CHF.

When combining the mean activations of all layers into a single global vector, the overall correlations (Fig. 14(b)) show the highest alignment with predicted and measured CHF, which substantiates that the internal representations of the model collectively encapsulate CHF-governing phenomena. Moreover, among the input parameters, heated length and outlet quality dominate the correlation. This indicates their

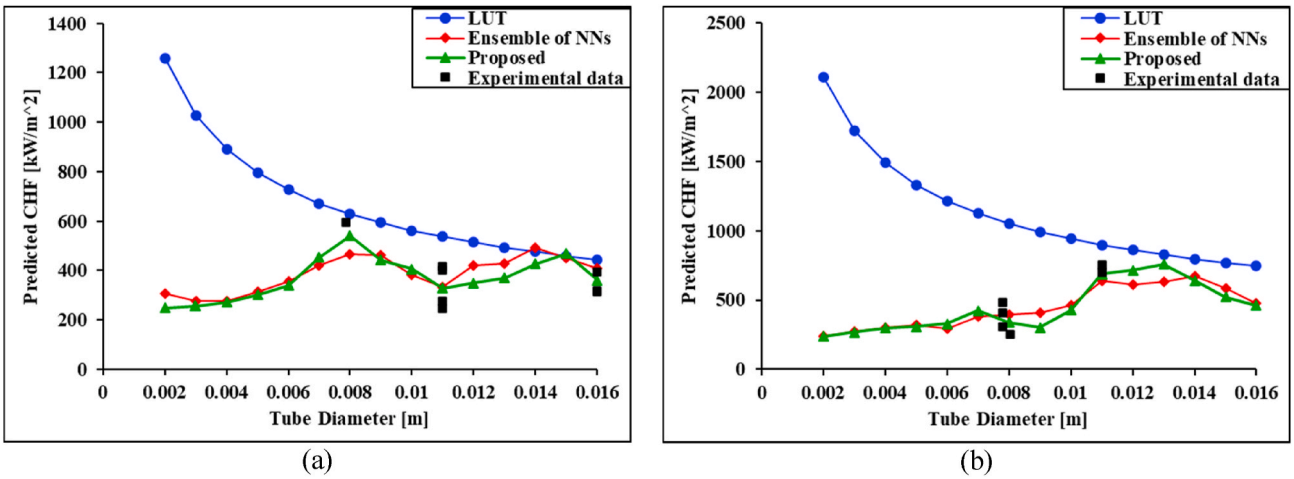

Fig. 15. Plots of physical relationships between CHF and diameter: (a) Slice #1 and (b) Slice #2.

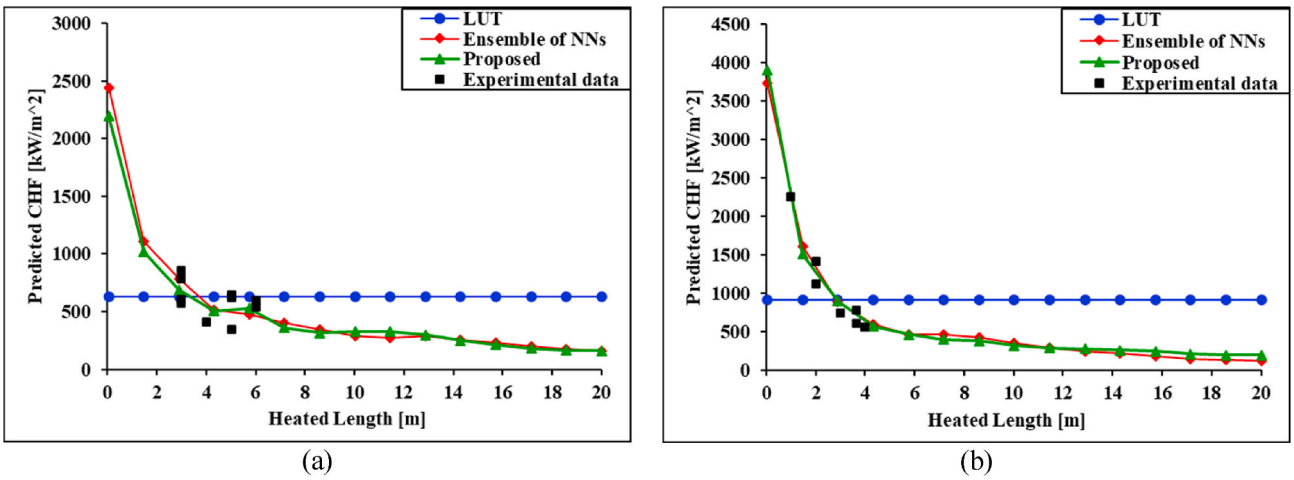

Fig. 16. Plots of physical relationships between CHF and heated length: (a) Slice #3 and (b) Slice #4.

influence on the global representation learned by the model. The negative correlation between the activation and the heated length is physically meaningful, as longer heated lengths generally reduce CHF due to increased heat input and associated vapor fraction. Similarly, the moderate negative correlation between the activation and the outlet quality reflects the tradeoff between vapor fraction and wall heat flux: high quality may increase risk of dryout, which, in turn, reduces CHF. The relatively low sensitivity to tube diameter, despite its geometric significance, may be attributed to the data normalization and the limited range of diameters in the dataset. This phenomenon has been reported in previous CHF modeling studies (Celata et al., 1996), where variations in CHF across tube diameters (e.g., 8 mm vs. 10–12 mm) were minimal,

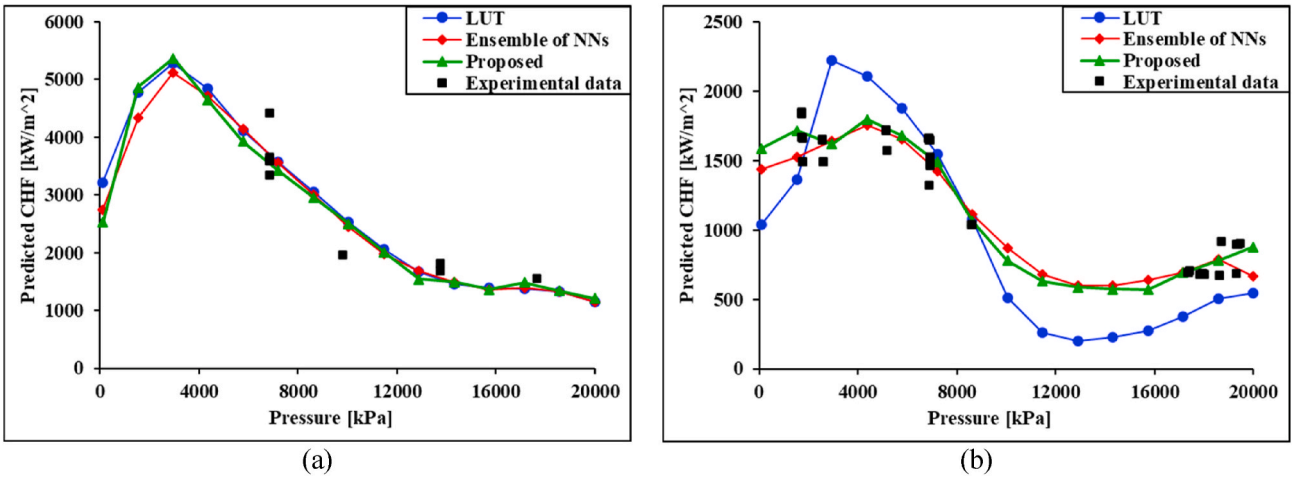

Fig. 17. Plots of physical relationships between CHF and pressure: (a) Slice #5 and (b) Slice #6.

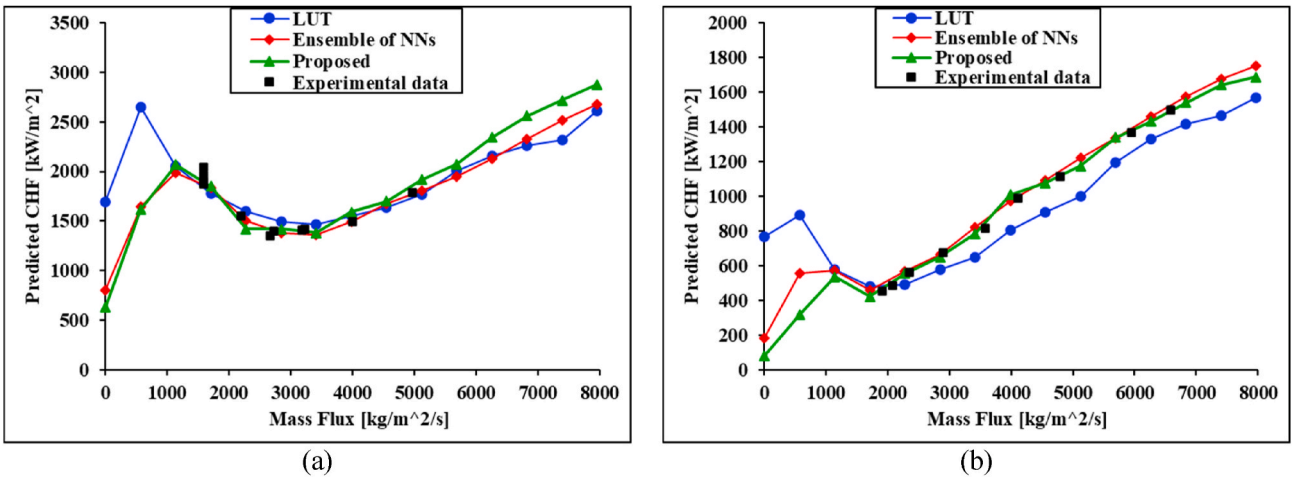

Fig. 18. Plots of physical relationships between CHF and mass flux: (a) Slice #7 and (b) Slice #8.

with CHF ratios ranging narrowly from approximately 0.96 to 1.02. This suggests that tube diameter exerts only a modest influence on CHF when its variation in the dataset is limited.

These findings reinforce the explainable nature of the proposed SA-1DCNN model, which show that the proposed model not only learns high-level abstractions that are aligned with target predictions, but it also maintains physical coherence with established thermohydraulic principles.

In conclusion, the practical applicability of the proposed explainable SA-1DCNN model extends beyond predictive accuracy to its potential role in reactor safety and regulatory processes. For licensing and safety evaluations, an interpretable prediction framework provides regulators and engineers with transparent evidence on how input parameters contribute to CHF limits, thereby strengthening the credibility of data-driven tools in safety assessments. In power uprate scenarios, where high thermal loads challenge the established CHF margins, the explainability of the model facilitates the identification of operating conditions most responsible for margin reductions, thereby offering a rational and aware basis for risk-informed decision-making. Similarly, for high burn-up operations, where uncertainties in fuel thermohydraulic behavior increase, interpretable feature attributions can confirm whether or not the learned representations of the model align with mechanistic CHF theories (e.g., hydrodynamic instability, dryout onset), thus reducing the likelihood of unrecognized extrapolation errors. By providing insights into both "what" the model predicts and "why", the proposed explainable SA-1DCNN model has the potential to complement mechanistic analyses and support reactor operators and regulators in making rational, aware, safety-oriented decisions.

5.3. Analysis of the physical behavior of the proposed SA-1DCNN model

The consistency of the predictions from the proposed SA-1DCNN method with the expected physical behavior is verified by using the slicing approach applied to the US NRC CHF database as presented in (Le Corre et al., 2024). In this approach, regions characterized by variations in one of the five key parameters have been recognized, while the remaining variables have been kept constant. From each of these regions, two representative slices were considered, as reported in Table 6. These "sliced" datasets enable the demonstration of the physical consistency of the proposed SA-1DCNN model predictions across the parameter space, thereby strengthening confidence in its predictive capability and mitigating the risk of overfitting. In conjunction with the "sliced" datasets, the corresponding experimental data are extracted based on the same variable ranges used in (Le Corre et al., 2024). The CHF prediction results for these ten "sliced" datasets are presented in Fig. 15 through Fig. 19 by comparing the proposed model, an optimized ensemble of NNs (Ahmed et al., 2025b) and the LUT. The ensemble of NNs is included as it represents a recently published CHF prediction model.

Fig. 15 presents the physical relationship between the CHF and channel diameter. The LUT, based on empirical correlations (Groeneveld et al., 2007), assumes that CHF decreases monotonically as

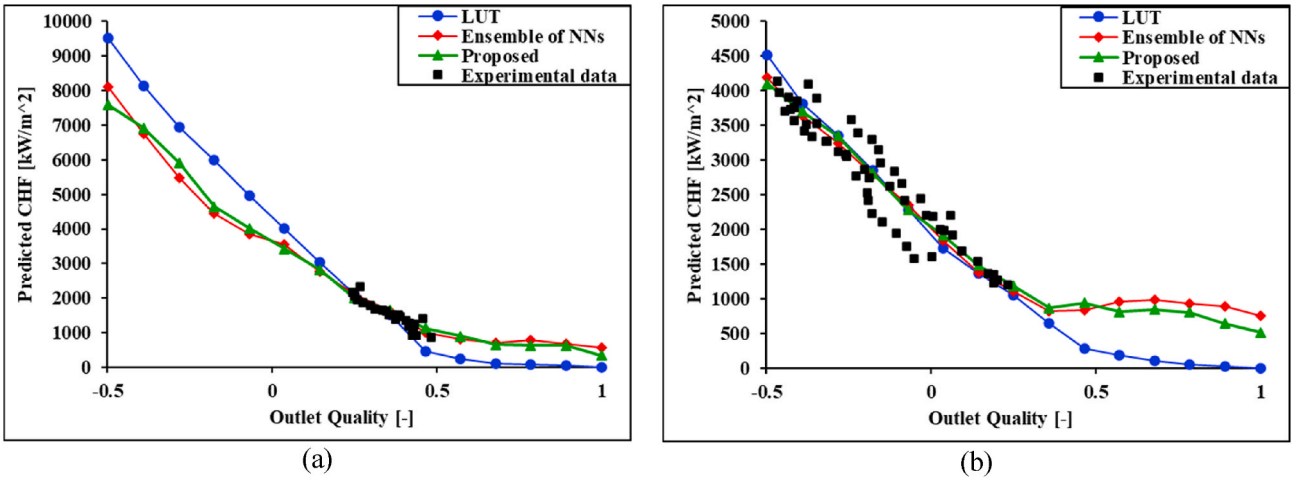

Fig. 19. Plots of physical relationships between CHF and outlet quality: (a) Slice #9 and (b) Slice #10.

diameter increases. This contradicts theoretical and experimental findings (Hall and Mudawar, 2000), which show that CHF actually increases with diameter under constant conditions. Both AI-based models correctly reflect this trend and align with experimental observations. In Fig. 16, the LUT model does not consider the effects of heated length, resulting in a constant trend. Contrary to LUT, the AI-based models show a physically consistent nonlinear behavior: CHF decreases up to approximately 10m, and then, it stabilizes, which reflects the dynamic behavior of heat flux distributions (Kinoshita et al., 2000). Fig. 17 shows that all three models demonstrate the nonlinear relationship between pressure and CHF, but the AI-based models are more closely align with experimental data. This alignment indicates variations in enthalpy-driven quality (Chun et al., 2001). In Fig. 18, CHF increases linearly with mass flux across all models, but the proposed model has the highest accuracy. Finally, Fig. 19 demonstrates that all models capture the inverse relationship between CHF and outlet quality, which aligns with vapor blanket formation (Yang et al., 2025). The proposed model provides the best match to experimental behavior.

6. Conclusions

In this paper, a new method for predicting CHF in vertical circular tubes using an explainable self-attention-enhanced one-dimensional CNN (SA-1DCNN) is proposed to improve predictive accuracy, generalization capabilities and interpretability. The proposed hybrid architecture combines convolutional layers for localized pattern extraction with self-attention mechanisms to capture global contextual dependencies among key physical parameters. Furthermore, Bayesian Optimization (BO) technique is employed to automatically set the optimal hyperparameters of the proposed SA-1DCNN method, which results in a robust and well-generalized model, optimized for CHF prediction. In addition to good predictive performance, this work emphasizes explainability and interpretability, achieved through comprehensive analyses, including mean activation profiling, spatial attention mapping and layer-wise correlation analysis. The proposed SA-1DCNN model is validated on the experimental CHF dataset provided by the WPRS-EGMUP task force on AI and ML project, a project supported by the OECD/NEA. In terms of the predictive performance on the CHF benchmark dataset, the proposed SA-1DCNN model significantly outperforms conventional and existing AI-based approaches across all metrics considered. Specifically, it reduces the MAPE by approximately 15 % and 71 % when compared to the best-performing (optimized ensemble of NNs) and worst-performing (LUT) comparison methods, respectively. Similarly, in terms of RMSPE, the proposed model shows a reduction of approximately 10 % and 71 % relative to the best-performing (optimized ensemble of NNs) and worst-performing (LUT) comparison methods, respectively. With respect to Q^2-error, the model achieves a reduction of approximately 21 % and 84 % when compared to the best-performing (Transformer) and worse-performing (LUT) comparison methods, respectively. Beyond its strong predictive capability, the interpretability and physical consistency of the model are shown through detailed explainability analyses. Here are the key findings.

- The average activation analysis of early layers of the proposed SA-1DCNN method shows that the convolutional layers extract evenly distributed local features with slightly higher sensitivity to pressure, whereas the attention layers focus on earlier spatial regions linked to geometric and thermodynamic parameters. This progression from local to global representation shows how the proposed model uses both convolutional and attention mechanisms to effectively capture nonlinear and interdependent patterns in CHF behavior.
- The group-wise activation analysis of the attention layers indicates that the proposed SA-1DCNN method emphasizes different parameters based on the CHF regime. In low CHF regimes, the attention focuses on pressure and vapor generation characteristics, whereas in high CHF regimes, geometric parameters are emphasized. This adaptable attention aligns with known thermophysical behavior and enhances both the interpretability and reliability of the model.
- The group-specific attention patterns in the SA-1DCNN model shift toward geometry for larger diameter, toward forcing at high pressure and low mass flux, and toward outlet quality at high values, reflecting classical CHF mechanisms: instability, dry-spot/microlayer and annular film dryout/entrainment, thus providing explainability-based validation of mechanistic consistency.
- The correlation analysis of layers activations shows that the model progresses from learning basic cues in early layers to capturing task-relevant representations in deeper layers. Mid-to-deep layers, particularly attention-enhanced ones and the final fully-connected dense layer, align strongly with CHF predictions, which reinforces the physical fidelity and explanatory ability of the model.
- When average activations of all layers are combined into a single vector, the correlations shows the highest alignment with the measured and predicted CHF. This indicates that the internal representations of the model collectively capture the phenomena that govern CHF. Among the input parameters, heated length and outlet quality dominate the correlations, which suggests that these parameters primarily influence the global representation learned by the model.
- In terms of global feature representation and physical alignment, the relatively low sensitivity to tube diameter is observed despite its geometric significance. This may be due to data normalization and the limited range of diameters in the dataset, a phenomenon that has been reported in the previous CHF modeling research works.
- The explainable SA-1DCNN model provides a transparent, mechanistically consistent basis for CHF prediction, thereby strengthening regulatory confidence and supporting safe decision-making in reactor licensing, power uprates and higher burn-up operations.

Finally, the analyses of physical behaviors of the proposed SA-1DCNN model further reaffirm the strong generalization capabilities of the SA-1DCNN model and its physical consistency with known thermohydraulic principles. These findings collectively underscore the effectiveness and trustworthiness of the SA-1DCNN model for CHF prediction in nuclear thermohydraulic systems.

The proposed SA-1DCNN model achieves strong predictive performance for CHF in circular water-cooled channels, effectively capturing nonlinear parameter interactions and offering interpretable learned representations. This highlights its potential as a supporting tool for reactor safety and performance analysis. Nevertheless, generalizability to non-circular geometries or alternative fluids is limited, given the dependence of CHF on flow geometry and fluid properties. Extension to such cases requires retraining with representative datasets or employing domain adaptation. Since the model is trained within the available parameter ranges, its extrapolation performance beyond these regimes is not guaranteed. Future research work will focus on: *1*) the use of intermediate layer activations for confidence estimation and uncertainty quantification to enable the model to assess whether a prediction should be trusted or rejected, an important step toward trustworthy deployment in nuclear thermohydraulic systems; and *2*) the incorporation of mechanistic constraints and physics-informed hybrid strategies to improve extrapolation robustness and extend applicability across channel types and fluids.

CRediT authorship contribution statement

Ibrahim Ahmed: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Enrico Zio:** Writing – review & editing, Visualization, Validation, Supervision, Resources, Methodology, Investigation, Formal analysis, 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.

Data availability

The authors do not have permission to share data.

References

Ahmed, I., Baraldi, P., Zio, E., Lewitschnig, H., 2025a. A data-driven modelling framework for predicting the quality of semiconductor devices to support burn-in decisions. Comput. Ind. Eng. 204, 111115. https://doi.org/10.1016/j.cie.2025.111115.

Ahmed, I., Gatti, I., Zio, E., 2025b. Optimized ensemble of neural networks for the prediction of critical heat flux. Nucl. Eng. Des. 439, 114111. https://doi.org/10.1016/j.nucengdes.2025.114111.

Biasi, L., Clerici, G.C., Tozzi, A., Sala, R., 1968. Extension of A.R.S. correlation to burnout prediction with non-uniform heating. J. Nucl. Energy 22, 705–716.

Bruder, M., Bloch, G., Sattelmayer, T., 2017. Critical heat flux in flow Boiling—Review of the Current understanding and experimental approaches. Heat Transf. Eng. 38, 347–360. https://doi.org/10.1080/01457632.2016.1189274.

Celata, G.P., Cumo, M., Katto, Y., Mariani, A., 1999. Prediction of the critical heat flux in water subcooled flow boiling using a new mechanistic approach. Int. J. Heat Mass Transf. 42, 1457–1466. https://doi.org/10.1016/S0017-9310(98)00286-5.

Celata, G.P., Cumo, M., Mariani, A., 1996. The effect of the tube diameter on the critical heat flux in subcooled flow boiling. Int. J. Heat Mass Transf. 39, 1755–1757. https://doi.org/10.1016/0017-9310(95)00264-2.

Chang, S., Baek, W.-P., 2003. Understanding, predicting, and enhancing critical heat flux. In: The 10th Int. Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-10). Seoul, Korea, p. October 5-9, 2003.

Chun, S.Y., Chung, H.J., Moon, S.K., Yang, S.K., Chung, M.K., Schoesse, T., Aritomi, M., 2001. Effect of pressure on critical heat flux in uniformly heated vertical annulus under low flow conditions. Nucl. Eng. Des. 203, 159–174. https://doi.org/10.1016/S0029-5493(00)00307-1.

Doroshchuk, V.E., Levitan, I.L., Lantzman, F.P., 1975. Investigation into Burnout in Uniformly Heated Tubes, vol. 75. ASME Publ.

Groeneveld, D.C., 2019. Critical Heat Flux Data Used to Generate the 2006 Groeneveld Lookup Tables. Nureg/Km-0011.

Groeneveld, D.C., Shan, J.Q., Vasić, A.Z., Leung, L.K.H., Durmayaz, A., Yang, J., Cheng, S.C., Tanase, A., 2007. The 2006 CHF look-up table. Nucl. Eng. Des. 237, 1909–1922. https://doi.org/10.1016/j.nucengdes.2007.02.014.

Hall, D.D., Mudawar, I., 2000. Critical heat flux (CHF) for water flow in tubes-I. Compilation and assessment of world CHF data. Int. J. Heat Mass Transf. 43, 2573–2604. https://doi.org/10.1016/S0017-9310(99)00191-X.

Han, J., Kamber, M., Pei, J., 2012. Data Mining: Concepts and Techniques, third ed. Morgan Kaufmann.

He, M., Lee, Y., 2018. Application of machine learning for prediction of critical heat flux: support vector machine for data-driven CHF look-up table construction based on sparingly distributed training data points. Nucl. Eng. Des. 338, 189–198. https://doi.org/10.1016/j.nucengdes.2018.08.005.

Hewitt, G.F., Hall-Taylor, N.S., 1970. Annular Two-phase Flow. Pergamon Press.

Ishii, M., Mishima, K., 1989. Droplet entrainment correlation in annular two-phase flow. Int. J. Heat Mass Transf. 32, 1835–1846. https://doi.org/10.1016/0017-9310(89)90155-5.

Ishii, M., Mishima, K., 1984. Two-fluid model and hydrodynamic constitutive relations. Nucl. Eng. Des. 82, 107–126. https://doi.org/10.1016/0029-5493(84)90207-3.

Jiang, B.T., Zhao, F.Y., 2013. Combination of support vector regression and artificial neural networks for prediction of critical heat flux. Int. J. Heat Mass Transf. 62, 481–494. https://doi.org/10.1016/j.ijheatmasstransfer.2013.03.025.

Katto, Y., 1992. A prediction model of subcooled water flow boiling CHF for pressure in the range 0.1-20 MPa. Int. J. Heat Mass Transf. 35, 1115–1123. https://doi.org/10.1016/0017-9310(92)90172-O.

Khalid, R.Z., Ahmed, I., Ullah, A., Zio, E., Khan, A., 2024. Enhancing accuracy of prediction of critical heat flux in circular channels by ensemble of deep sparse autoencoders and deep neural Networks. Nucl. Eng. Des. 429, 113587.

Khalid, R.Z., Ullah, A., Khan, Asifullah, Khan, Afrasyab, Inayat, M.H., 2023. Comparison of standalone and hybrid machine learning models for prediction of critical heat flux in vertical tubes. Energies 16, 3182. https://doi.org/10.3390/en16073182.

Kingma, D.P., Ba, J.L., 2015. Adam: a method for stochastic optimization. In: 3rd Int. Conf. Learn. Represent. ICLR 2015 - Conf. Track Proc, pp. 1–15.

Kinoshita, H., Yoshida, T., Nariai, H., Inasaka, F., 2000. Effect of heated length on the critical heat flux of subcooled flow boiling. Part 1: observation of bubbles and slug length at atmospheric pressure. Heat Tran. Res. 29, 132–143. https://doi.org/10.1002/(sici)1523-1496(200003)29:2<132::aid-htj4>3.0.co;2-t.

Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M., Inman, D.J., 2021. 1D convolutional neural networks and applications: a survey. Mech. Syst. Signal Process. 151, 107398. https://doi.org/10.1016/j.ymssp.2020.107398.

Le Corre, J.-M., Delipei, G., Wu, X., Zhao, X., 2024. Benchmark on artificial intelligence and machine learning for scientific computing in nuclear engineering. Phase 1: critical heat flux exercise specifications. NEA Working Papers.

Lecun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. Nature 521, 436–444. https://doi.org/10.1038/nature14539.

LeCun, Y.A., Bottou, L., Orr, Genevieve B., Müller, K.-R., 2012. Efficient BackProp. In: Montavon, G., Orr, Geneviève B., Müller, K.-R. (Eds.), Neural Networks: Tricks of the Trade. Springer, Berlin, Heidelberg, pp. 9–48. https://doi.org/10.1007/978-3-642-35289-8_3.

Lee, C.H., Mudawwar, I., 1988. A mechanistic critical heat flux model for subcooled flow boiling based on local bulk flow conditions. Int. J. Multiphas. Flow 14, 711–728. https://doi.org/10.1016/0301-9322(88)90070-5.

Lee, M.R., Jeong, H.J., Choi, Y.J., Gatton, T.M., 2006. A nuclear power plant expert system using artificial neural networks. In: Ladvances in Neural Networks - ISNN 2006. ISNN 2006. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, pp. 1239–1245. https://doi.org/10.1007/11760023_180.

Mao, C., Jin, Y., 2024. Uncertainty quantification study of the physics-informed machine learning models for critical heat flux prediction. Prog. Nucl. Energy 170, 105097. https://doi.org/10.1016/j.pnucene.2024.105097.

Mazzola, A., 1997. Integrating artificial neural networks and empirical correlations for the prediction of water-subcooled critical heat flux. Rev. Gen. Therm. 36, 799–806. https://doi.org/10.1016/S0035-3159(97)87750-1.

Močkus, J., 1974. On Bayesian methods for seeking the extremum. In: Marchuk, G.I. (Ed.), Optimization Techniques IFIP Technical Conference. Springer, Berlin, Novosibirsk, pp. 400–404. https://doi.org/10.1007/3-540-07165-2_55.

Morse, R.W., Moreira, T.A., Chan, J., Dressler, K.M., Ribatski, G., Hurlburt, E.T., McCarroll, L.L., Nellis, G.F., Berson, A., 2021. Critical heat flux and the dryout of liquid film in vertical two-phase annular flow. Int. J. Heat Mass Transf. 177, 121487. https://doi.org/10.1016/j.ijheatmasstransfer.2021.121487.

Myers, P.A., Panczyk, N., Chidige, S., Craig, C., Cooper, J., Joynt, V., Radaideh, M.I., 2025. pyMAISE: a Python platform for automatic machine learning and accelerated development for nuclear power applications. Prog. Nucl. Energy 180, 105568. https://doi.org/10.1016/j.pnucene.2024.105568.

OECD/NEA, 2022. Task force on artificial intelligence and machine learning for scientific computing in nuclear engineering [WWW Document]. URL. https://www.oecd-nea.org/jcms/pl_77779/task-force-on-artificial-intelligence-and-machine-learning-for-scientific-computing-in-nuclear-engineering. accessed 5.5.25.

Osborne, J.W., 2010. Improving your data transformations: applying the Box-Cox transformation. Pract. Assessment, Res. Eval. 15, 1–9.

Qi, S., Han, B., Zhu, X., Yang, B.W., Xing, T., Liu, A., Liu, S., 2025. Machine learning in critical heat flux studies in nuclear systems: a detailed review. Prog. Nucl. Energy 179, 105535. https://doi.org/10.1016/j.pnucene.2024.105535.

Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning representations by back-propagating errors. Nature 323, 533–536.

Saha, P., Zuber, N., 1974. Point of net vapor generation and vapor void fraction in subcooled boiling. In: International Heat Transfer Conference. Scripta Book Co., Tokyo, Japan https://doi.org/10.1615/ihtc5.430.

Sajjad, U., Mehdi, S., Hussain, I., Rehman, T. ur, Sultan, M., Rashidi, M.M., Yan, W.M., 2025. Physics driven interpretable deep learning-based insights into boiling crisis of smooth and roughened surfaces. Alex. Eng. J. 116, 112–128. https://doi.org/10.1016/j.aej.2024.10.103.

Stano, M., Benesova, W., Martak, L.S., 2020. Explaining Predictions of Deep Neural Classifier via Activation Analysis.

Todreas, N.E., Kazimi, M., 2011. Nuclear Systems Volume I: Thermal Hydraulic Fundamentals, second ed. CRC Press.

Tong, L.S., 1967. Prediction of departure from nucleate boiling for an axially non-uniform heat flux distribution. J. Nucl. Energy 21, 241–248. https://doi.org/10.1016/S0022-3107(67)90054-8.

Tong, L.S., Tang, Y.S., 1997. Boiling Heat Transfer and Two-phase Flow, first ed. Routledge. https://doi.org/10.1201/9781315138510.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. Adv. Neural Inf. Process. Syst. 2017-Decem 5999–6009.

Wang, K., Wang, D., Liu, X., Cheng, S., Wang, S., Zhou, W., Miwa, S., Okamoto, K., 2025. Re-examining the input-parameters and AI strategies for Critical Heat Flux prediction. Energy 318, 134606. https://doi.org/10.1016/j.energy.2025.134606.

Wang, L., Dernoncourt, F., Bui, T., 2020. Bayesian optimization for selecting efficient machine learning models. arXiv:2008.00386 [cs.LG]. https://doi.org/10.48550/arXiv.2008.00386.

Weisman, J., Pei, B.S., 1983. Prediction of critical heat flux in flow boiling at low qualities. Int. J. Heat Mass Transf. 26, 1463–1477.

Yang, J., Yong, H., Kim, S., Park, I.W., Lee, Y.G., Vadlamudi, S.R.G., Park, H.S., 2025. Critical heat flux dependence on surface orientation and bubble dynamics in pool boiling over silicon and silicon dioxide surfaces. Int. J. Energy Res., 6413134 https://doi.org/10.1155/er/6413134, 2025.

Zhang, Q., Zhu, 2018. Visual interpretability for deep learning: a survey. Front. Inf. Technol. Electron. Eng. 19, 27–39. https://doi.org/10.1631/FITEE.1700808.

Zhao, X., Shirvan, K., Salko, R.K., Guo, F., 2020. On the prediction of critical heat flux using a physics-informed machine learning-aided framework. Appl. Therm. Eng. 164, 114540. https://doi.org/10.1016/j.applthermaleng.2019.114540.

Zhou, W., Miwa, S., Wang, H., Okamoto, K., 2024. Assessment of the state-of-the-art AI methods for critical heat flux prediction. Int. Commun. Heat Mass Tran. 158, 107844. https://doi.org/10.1016/j.icheatmasstransfer.2024.107844.

Zuber, N., 1959. Hydrodynamic Aspects of Boiling Heat Transfer. Univ. of California, Los Angeles, CA (United States). https://doi.org/10.2172/4175511.