Computational Materials Science 274 (2026) 115025

Contents lists available at ScienceDirect

Computational Materials Science

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

Full length article

Towards real-time prediction of thermal history and hardness in laser powder bed fusion using deep learning

Philipp Schüßler *, Volker Schulze, Stefan Dietrich

Karlsruhe Institute of Technology (KIT), Institute for Applied Materials - Materials Science and Engineering (IAM-WK), Karlsruhe, 76131, Germany

ARTICLE INFO

Keywords:
Additive manufacturing
Deep learning
Surrogate model
Carbon steel
AISI 4140
42CrMo4
Thermal history
Hardness

ABSTRACT

Laser powder bed fusion (PBF-LB) of quenched-and-tempered steels is governed by highly localized thermal histories that control microstructure evolution and hardness. Predicting these process–structure–property relationships typically requires computationally intensive finite element (FE) simulations, limiting real-time applicability. Here, we present a physics-informed deep learning surrogate for real-time prediction of thermal histories and hardness in PBF-LB of AISI 4140. The model integrates process parameters and spatial laser–material interactions within an autoregressive sequence framework to capture path-dependent thermal behavior. Trained on multiscale FE data, the model reconstructs local temperature–time histories with high fidelity and enables hardness prediction via a non-isothermal Hollomon–Jaffe relationship. A two-layer LSTM ensemble achieves a temperature RMSE of (2.6 ± 1.2) K on an independent test set. The central contribution is the prediction of complete local thermal histories and their propagation through a validated tempering model to obtain local hardness under varying process conditions and cross-section geometries. Beyond this specific system, the proposed framework establishes a generalizable strategy for bridging high-fidelity simulation and real-time prediction in manufacturing processes. It thus provides a foundation for accelerated process design, in situ decision-making, and the realization of digital twins in materials processing.

1. Introduction

Laser Powder Bed Fusion (PBF-LB) offers high geometric design freedom and enables the fabrication of complex, functionally integrated components. Despite this potential, industrial use is still largely limited to well-established alloys, while high-strength martensitic quenched-and-tempered steels such as AISI 4140 (42CrMo4) remain challenging due to cracking, residual stresses and difficult process control [1].

In these steels the local microstructure and hardness are governed by the intrinsic heat treatment during the PBF-LB process. This treatment takes the form of repeated, highly localized heating and cooling cycles during layerwise exposure [2]. The resulting process–structure–property relationships are strongly driven by the local thermal history and are difficult to characterize experimentally, especially under the extreme heating and cooling rates typical of PBF-LB [1–3]. This motivates the development of efficient, predictive models for local temperature histories and derived properties as a basis for process optimization and, ultimately, real-time data-driven process control.

Numerical simulations can provide detailed insight into the intrinsic transformation and tempering phenomena, help identify their governing causes and quantify their impact on microstructure and mechanical properties [1]. In particular, finite element (FE) models have become an important tool for predicting temperature evolution, residual stresses and microstructure in PBF-LB [4–6]. Most existing FE studies, however, focus on alloys without solid-state phase transformations (e.g. 316L or aluminum alloys) and employ simplified macroscopic or mesoscale models with effective material properties and averaged energy input [3,7,8]. In martensitic steels with complex transformation and tempering behavior, however, these simplified energy input representations fail to capture the highly localized heating and cooling rates. These localized rates control phase fraction evolution and ultimately determine the resulting hardness [9]. Moreover, their high computational cost limits their applicability for real-time prediction and rapid process exploration.

The underlying challenge is schematically illustrated in Fig. 1, which highlights the localized and path-dependent thermal histories in PBF-LB, their governing role in process-structure–property relationships, and the limitations of conventional FE-based approaches for efficient process exploration.

* Corresponding author.
E-mail address: philipp.schuessler@kit.edu (P. Schüßler).

https://doi.org/10.1016/j.commatsci.2026.115025
Received 15 April 2026; Received in revised form 24 July 2026; Accepted 18 August 2026
Available online 28 August 2026

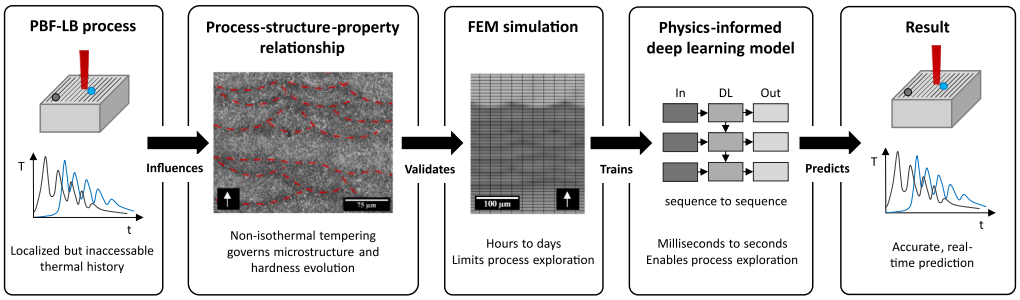

Fig. 1. Conceptual overview of process-structure–property relationships in PBF-LB and the associated computational challenge. Localized, path-dependent thermal histories govern microstructure evolution and resulting hardness, but their accurate prediction via multiscale FE simulations is computationally expensive. A physics-informed deep learning model enables efficient prediction of thermal histories and resulting properties, facilitating rapid process exploration.

2. Literature review and research gap

In parallel, the attention to data-driven methods and deep learning (DL) in additive manufacturing has increased in recent years [10–12]. A wide range of neural network architectures have been applied to tasks such as process monitoring, defect detection, process parameter optimization and prediction of final part properties. Physics-informed machine learning (PIML) in additive manufacturing comprises several distinct modes of incorporating physical knowledge. These include physics-based feature engineering, physics-based shaping of the model architecture and physics-based modification of the loss function [13]. This distinction is important because models trained on high-fidelity simulation data with engineered physical descriptors differ conceptually from physics-informed neural networks (PINNs), in which governing equations and boundary or initial conditions are imposed through the training objective.

For time-dependent thermal processes, recurrent neural networks (RNNs) and their gated variants are natural candidates for representing path-dependent temperature sequences [14–16]. Mozaffar et al. used a Gated Recurrent Unit (GRU)-based RNN to predict pointwise thermal histories in Directed Energy Deposition (DED) across variations in geometry, toolpath and process parameters [17]. Ren et al. subsequently combined an RNN with a deep neural network and trained the model on data from an experimentally validated numerical simulation to map laser scanning patterns to thermal-history distributions [18]. A complementary non-neural approach was presented by Ness et al. who used extremely randomized trees and generic, physics-engineered features derived from FE simulations to obtain transferable nodal thermal-history predictions across several geometries and deposition conditions [19]. These studies demonstrate the effectiveness of temporal state representations and physics-derived inputs, but predominantly concern DED or simplified material descriptions without a coupled solid-state transformation and hardness model.

Alternative architectures emphasize different aspects of the thermal prediction problem. Recurrent graph neural networks (RGNNs) explicitly retain mesh connectivity and have been used to predict long thermal histories for geometries not included in training [20]. More recently, transfer learning has enabled RGNN temperature prediction across new DED geometries, process conditions and material systems [21]. Temporal convolutional networks (TCNs) provide a convolution-based sequence alternative and achieved accuracy comparable to previously proposed deep-learning thermal surrogates at a fraction of their compute and training times in DED [22]. For PBF-LB, the STEAM framework compared LSTM and GRU variants using temporal and spatial physics-based features for part-scale scanwise thermal histories and demonstrated transfer to unseen, larger geometries [23]. Consequently, neither thermal-history prediction nor transfer across geometry alone constitutes the research gap addressed here.

PINNs follow a different formulation by incorporating the governing equations directly into learning. Zhu et al. embedded conservation of mass, momentum and energy to predict temperature and melt-pool fluid dynamics in metal additive manufacturing [24]. Hosseini et al. obtained parametric transient and steady-state temperature profiles and melt-pool dimensions for single-track PBF-LB through an unsupervised PINN formulation [25]. Peng and Panesar extended PINN-based thermal-history prediction to multilayer DED and specifically addressed discontinuities introduced by layerwise activation [26]. These works show the potential of equation-constrained learning, while also illustrating the additional formulation and training challenges associated with moving heat sources, layer activation and material nonlinearities.

A separate body of research has linked process or thermal information to microstructure and properties. Kuehne et al. predicted melt-track geometry and part density in PBF-LB from process parameters using classical machine-learning algorithms [27]. Gan et al. combined thermal-fluid simulations, experimental measurements, mechanistic relations and a self-organizing map to design microstructure and microhardness in deposited Ni-based superalloys [28], while Wolff et al. experimentally validated a computational thermal-history–microhardness relationship for laser-deposited Inconel 718 on carbon steel [29]. Xie et al. transformed measured DED thermal histories into wavelet representations and used a convolutional neural network to predict location-dependent as-built mechanical properties [30]. For PBF-LB, Yu et al. predicted Ti–6Al–4V hardness directly from process parameters using explainable machine-learning models [31]. These studies establish important process–property relationships, but they either concern other AM processes and alloys or predict final-state quantities without first reconstructing the complete local multicycle thermal history.

The suitability of a model family therefore depends on the representation and prediction target. CNNs are advantageous for regular spatial fields or transformed signal representations, GNNs for explicitly connected full-field meshes, TCNs for convolutional sequence modeling and PINNs for imposing governing equations. The present task instead consists of pointwise, autoregressive sequences containing short laser-induced heating events, longer cooling intervals and dependencies over the complete layer-processing sequence. An LSTM is therefore a task-appropriate choice for the present process–property formulation.

Taken together, the literature has separately demonstrated fast thermal-history reconstruction, geometry-transferable thermal surrogates and data-driven prediction of hardness or other mechanical properties. To the best of the authors' knowledge, the remaining gap is the integration of these elements in a computationally efficient process–structure–property framework that predicts complete local multicycle temperature histories for PBF-LB of a martensitic quenched-and-tempered steel and propagates the histories through transformation and non-isothermal tempering models to obtain local hardness across process and cross-section variations.

In the present work, the term physics-informed therefore refers to training data generated by an experimentally validated FE model, physics-derived input features and coupling of the predicted thermal history to a physically based tempering relationship. The governing

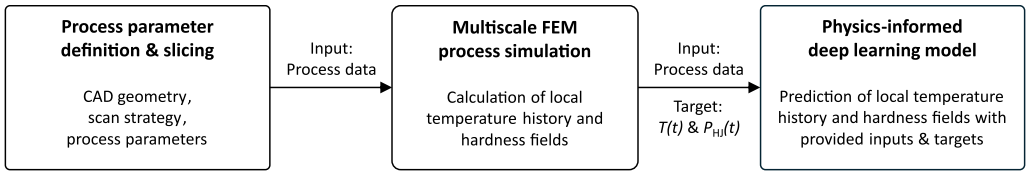

Fig. 2. Schematic overview of the data flow used for training. Process definition and slicing provide scan vectors and process parameters, which serve as input to a multiscale FEM process model that computes local temperature–time histories and derived tempering quantities (Hollomon-Jaffe parameter P_{HJ} and hardness HV1). The FEM simulation results are used to train a physics-informed sequence-to-sequence LSTM ensemble, which subsequently acts as a fast surrogate model, enabling accelerated process optimization.

heat equation is not imposed directly in the loss function. The approach is primarily based on physics-informed feature engineering and simulation-generated data and is distinct from a classical PINN [13].

The present study addresses this gap by developing and analyzing a deep learning based surrogate model for PBF-LB of quenched-and-tempered steels with a solid phase transformation like AISI 4140. The study focuses on a physics-informed LSTM surrogate and examines model-capacity sensitivity within this recurrent architecture. The overall training and inference workflow is illustrated in Fig. 2.

High-fidelity multiscale FE simulations are used to generate local thermal histories and derive hardness via the Hollomon–Jaffe tempering parameter, which serve as training data for a physics-informed sequence-to-sequence LSTM model [32–34]. The overarching goal is to enable fast and accurate prediction of local temperature–time histories and resulting hardness for a broad set of process and geometric conditions, thereby providing a basis for process optimization and, in the longer term, for real-time in-process property control.

More specifically, the objectives of this work are

- to demonstrate the feasibility of a physics-informed LSTM surrogate for predicting local temperature histories from process and geometric features, and to examine the sensitivity of its performance to LSTM model capacity,
- to evaluate the predictive performance of the LSTM ensemble relative to FE simulations using global (RMSE) and process-specific metrics (maximum temperature, final temperature and hardness),
- to analyze feature relevance and assess model generalization across unseen process conditions and geometries, and
- to quantify the computational speed-up of the surrogate model and its potential for accelerated and real-time process optimization.

By combining a physics-based multiscale FE model with a carefully designed LSTM surrogate, this work aims to contribute to a physically grounded, yet computationally efficient, framework for modeling the process–structure–property relationships in PBF-LB of martensitic quenched-and-tempered steels with a clear pathway toward real-time prediction. The best-performing of three investigated LSTM capacities is used as the reference configuration for the subsequent analyses.

3. Methods

3.1. Finite element process simulation

A finite element (FE) model was employed to generate the physically consistent reference data used for training and validating the deep learning model. The FE framework was implemented in Abaqus/Standard and extended by user-defined subroutines written in Fortran. It was based on a modular, Python-controlled workflow consisting of four primary modules: (1) scan strategy analysis, (2) pre-processing, (3) FE simulation, and (4) post-processing. This modular structure enabled automated transfer of process and material parameters, boundary conditions, and geometric data between modules, ensuring reproducibility across process configurations. A detailed experimental validation was previously published by the authors [35]. The process configurations underlying the validated FE model were based on build jobs manufactured on an SLM280 system (SLM Solutions Group AG, Germany) equipped with a Yb:YAG fiber laser with a maximum power of 400 W, a wavelength of 1070 nm and a nominal focal diameter of 80 µm. Processing was conducted under argon using a layer thickness of 30 µm and a bidirectional hatch strategy.

The simulation described transient heat conduction during the Laser Powder Bed Fusion (PBF-LB) process. The laser energy input was modeled using a double-ellipsoidal Goldak heat source, which was calibrated against experimentally measured melt pool geometries according to Hocine et al. [6] and Zhang et al. [36]. Temperature- and phase-dependent thermophysical properties of the low alloyed quench and tempering steel AISI 4140 (42CrMo4) were implemented based on experimentally fitted material models derived from literature and the authors previous publications [35]. The material models included solid–liquid transitions [35], austenite transformation [37], bainitic transformation [38,39], martensitic transformation [35,40], and carbide precipitation (ε-carbide and cementite) during tempering [9,41]. All phase transformations were modeled using temperature-dependent modified Johnson–Mehl–Avrami–Kolmogorov (JMAK) equations, implemented via `UMATHT` user subroutines. A modified Hollomon-Jaffe equation for non-isothermal tempering was used in order to calculate the resulting hardness distribution based on the thermal history and the phase transformations [33,35]. Thermal boundary conditions were defined as follows: the bottom surface of the build plate was maintained at a fixed base plate temperature (BPT), as measured on the machine. The lateral surfaces were assumed adiabatic due to the low thermal conductivity of the surrounding powder bed and the top surface was subjected to combined convective and radiative heat losses, modeled using a constant heat transfer coefficient $h = 20\,\mathrm{W\,m^{-2}\,K^{-1}}$ and an emissivity $\varepsilon = 0.5$, following Peyre et al. [42].

Each scan vector was characterized by its start and end coordinates $(x_{\mathrm{start}}, y_{\mathrm{start}}, z_{\mathrm{start}})$ and $(x_{\mathrm{end}}, y_{\mathrm{end}}, z_{\mathrm{end}})$, laser power P, scanning speed v, hatch angle α_{SV}, and layer thickness t_S. The process parameters were automatically extracted from the machine control files (`.3mf`, `.lsr`) using the Python-based preprocessor module (Project "3mf-to-pbflb-gcode" [43]). For each simulated layer, the laser motion was discretized into sequential time steps corresponding to active and inactive laser phases. The applied multiscale approach combined Goldak-based high-resolution near-field modeling with scan-vector- and area-equivalent laser source formulations to reduce computational cost while maintaining physical accuracy for the far-field [35]. The multiscale approach is visualized in Fig. 3. The simulation output included geometrically and temporally resolved temperature fields $T(x, y, z, t)$, local phase fractions, the Hollomon–Jaffe tempering parameter P_{HJ} [32,33], and the derived Vickers hardness HV. The hardness–P_{HJ} correlation was implemented according to the validated relationship for AISI 4140 by Schüßler et al. [35].

3.2. Deep learning model

The prediction of the temporal temperature evolution was based on a recurrent neural-network architecture, specifically a sequence-to-sequence Long Short-Term Memory (LSTM) network. An LSTM is a

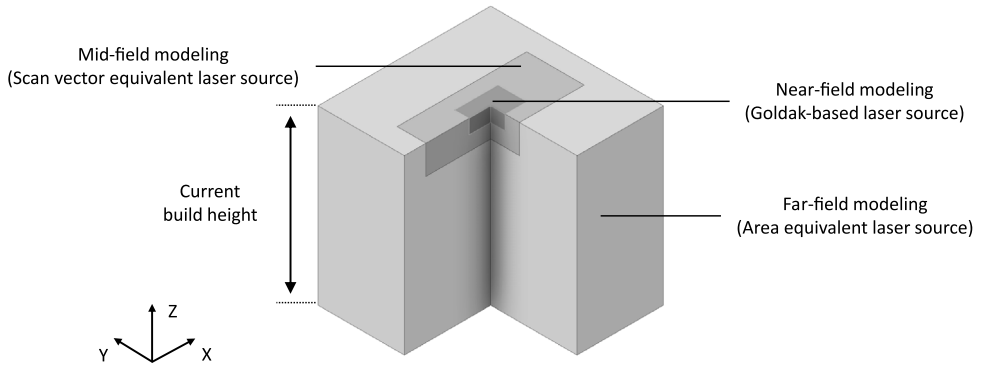

Fig. 3. Visualization of the multiscale modeling approach for the laser energy input of the part at a specific build height. Near-field Goldak-based laser source in the center (dark gray), a scan vector equivalent laser source for the mid-field (gray) and an area-equivalent laser source for the far-field modeling (light gray).

gated form of RNN whose hidden and cell states are designed to retain long-range temporal information and mitigate the vanishing-gradient limitations of conventional RNNs [14,44]. This is suitable for the present pointwise sequences, which contain short laser-induced heating events, longer cooling intervals and dependencies over the complete layer-processing sequence. The autoregressive formulation also reflects the temporal state evolution by using the preceding temperature in the prediction of the subsequent time step. Temporal convolutional networks constitute a credible convolution-based alternative for thermal sequences [22], whereas CNNs applied to transformed thermal signals and GNNs applied to explicitly connected meshes can exploit regular image-like representations or spatial connectivity, respectively [20,21,30]. In the present dataset, spatial information is represented by physics-derived laser–point distances and process descriptors rather than a full-field mesh. The selected LSTM formulation is therefore aligned with the pointwise autoregressive target. All models were implemented in PyTorch using the `nn.Module` class.

3.2.1. Dataset preparation

The dataset used for training and validation of the LSTM models was generated with the multiscale finite element (FE) process model described in Section 3.1. Only evaluation points from the mesoscale domain with the implemented Goldak heat source were considered, as these regions exhibit the most accurate temperature histories for the investigated process parameters. The geometries were fabricated directly on the build plate without support structures. Due to the simultaneous fabrication of multiple samples per build job, each layer cooled completely to the prescribed base plate temperature (BPT) before the next exposure. This dwell time was included in the simulated thermal histories.

The time-dependent FE results were exported in `.npy` format using a dedicated post-processing routine with the Abaqus/Standard Python-API and subsequently converted to compressed `HDF5` files for efficient storage and handling. Each record contains the process parameters (P, v, d_{Fokus}, h, α_{SV}, BPT), the scan vector geometry (start and end coordinates in x, y, z), the coordinates of the evaluation point, and the corresponding time-dependent outputs (temperature $T(t)$, phase fractions $w(t)$, Hollomon–Jaffe parameter $P_{\mathrm{HJ}}(t)$, and hardness $HV1$).

Simulations were carried out for three process parameter combinations (P006, P030, P076), three BPTs (non-controlled, 100 °C, 200 °C), and four geometries with different scan vector lengths (3 mm, 6 mm, 9 mm and 12 mm). The full dataset comprises approximately 650 000 evaluation points with a compressed size of about 250 GB. For LSTM training, a subset of 32 306 points was randomly selected and split into training (80 %), validation (10 %), and test (10 %) sets, ensuring a strict separation between these subsets to avoid data leakage.

The feature selection followed a correlation analysis to remove redundant inputs. The final feature set included: laser power P, laser on/off state, scan speed v, cyclically transformed hatch angle components $\sin(\alpha_{\mathrm{SV}})$ and $\cos(\alpha_{\mathrm{SV}})$, relative distances between laser and evaluation point in x-, y-, and z-direction (scan-vector coordinate system), vertical distance between build plate and evaluation point in the global z-direction, time since the end of laser exposure, BPT, and the temperature at the previous time step T_{t-1} (autoregressive feature). All input variables were normalized using min–max scaling based on the range of the training data. The target variable for all models was the complete temperature sequence $T(t)$ at each evaluation point.

Table 1
LSTM architectures investigated in this study.

Model	LSTM layers	Hidden units per layer	Trainable parameters
A	1	16	1937
B	2	32	14 369
C	4	64	119 873

3.2.2. Network architecture

To predict the full temperature history, a sequence-to-sequence LSTM architecture was employed. At each time step t, the network receives the current input features, the internal hidden and cell states and outputs the temperature at the same time step. This operation is repeated over all time steps with a fixed time increment of 2×10^{-4} s. A simplified version of the LSTM was used for a previous small scale study by the authors [34]. Three LSTM architectures with increasing capacity were evaluated to assess sensitivity to model size within the selected recurrent model family (see Table 1).

Each architecture consists of stacked LSTM layers followed by a fully connected output layer that maps the final hidden state at each time step to the predicted temperature. During training, a teacher-forcing strategy [45] was used, in which the reference temperature from the previous time step is provided as auto-regressive input to stabilize learning and improve convergence. During validation, testing and inference, the model operates in an auto-regressive mode, recursively feeding its own predictions back as input for the subsequent time step.

3.2.3. Training strategy

The network parameters were optimized by minimizing the mean squared error (MSE) between predicted and reference temperatures. The Adam optimizer was employed with an adaptive learning-rate scheduler. To promote robust convergence and generalization, a curriculum learning strategy was implemented: the network was first trained on simpler examples with shorter sequences and smaller cross sections and then gradually exposed, in six stages, to increasingly complex data with larger cross sections (more scan vectors per layer) and longer sequences. Early stopping based on the validation loss was used to prevent overfitting. For each architecture (Models A–C), an

ensemble of five networks was trained from different random initializations. The ensemble mean was used as the final prediction, while the standard deviation across ensemble members provided a measure of predictive uncertainty. All training and inference runs were performed on a workstation equipped with an Nvidia RTX 4090 GPU. Inference for multiple evaluation points was parallelized over batches for efficiency.

3.2.4. Numerical validation

The accuracy of the predicted temperature sequences $\hat{T}(t)$ was evaluated by comparison with the FE reference data $T_{\mathrm{FEM}}(t)$ on the test dataset. The root mean square error (RMSE) was used as the primary metric and defined as

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\hat{T}_i - T_{\mathrm{FEM},i}\right)^2}, \tag{1}$$

where N denotes the number of time steps in the sequence.

In addition, several process-specific error metrics were computed for each evaluation point: (i) the error of the maximum temperature,

$$\Delta T_{\max} = \max_t \hat{T}(t) - \max_t T_{\mathrm{FEM}}(t), \tag{2}$$

(ii) the temperature error at the final time step,

$$\Delta T_{\mathrm{end}} = \hat{T}(t_{\mathrm{end}}) - T_{\mathrm{FEM}}(t_{\mathrm{end}}), \tag{3}$$

and (iii) the error of the resulting hardness,

$$\Delta HV = \hat{HV} - HV_{\mathrm{FEM}}, \tag{4}$$

where $\hat{HV}$ and HV_{FEM} are obtained from the predicted and FE-based temperature histories via the Hollomon–Jaffe parameter. These metrics provide complementary information on amplitude accuracy and temporal alignment of predicted and simulated temperature histories.

3.3. Model analysis and generalization framework

3.3.1. Permutation feature importance

The relative importance of the individual input features was quantified using Permutation Feature Importance (PFI). For each feature, its values were randomly permuted over the test dataset while all other inputs were kept unchanged. The resulting change in RMSE, averaged over multiple permutations, was interpreted as a measure of the feature's influence on the model predictions (Fig. 6). For the autoregressive temperature input T_{t-1}, simple permutation would destroy the temporal structure of the sequence. Instead, T_{t-1} was replaced by a constant value and the resulting RMSE was compared to the baseline.

3.3.2. Error analysis across process categories

To investigate systematic dependencies of the prediction error on specific process conditions, the RMSE values for the test dataset were analyzed as a function of different process categories: (i) process parameter combination (P006, P030, P076), (ii) BPT (non-controlled, 100 °C, 200 °C), (iii) scan vector length (and thus cross-sectional area), and (iv) distance of the evaluation point from the current part surface. For each category, the distribution of RMSE values was visualized using boxplots (Fig. 7), providing insight into the median error, spread and occurrence of outliers under different process conditions.

3.3.3. Generalization tests for process parameters and geometry

The generalization capability of Model B was further assessed by applying the trained network to input conditions outside the discrete combinations used for training. For process parameters, new test datasets were generated in which hatch angles, BPT, laser power P and scan speed v were systematically varied, both individually and in combination. The corresponding error metrics (RMSE, $\Delta T_{\max}$, ΔT_{end}, ΔHV) were evaluated and summarized in Table 4, and representative temperature–time histories were compared with FE simulations (Fig. 8).

Table 2
Error metrics of the three LSTM architectures for training, validation, and test datasets. Values are reported as mean ± standard deviation. The training data used is described in Section 3.2.1.

Model	Dataset	RMSE/K	$\Delta T_{\max}$/K	ΔT_{end}/K	ΔHV/HV1
A	Train	4.4 ± 2.0	55 ± 105	4.2 ± 3.2	10 ± 28
	Val	4.4 ± 2.0	54 ± 102	4.1 ± 3.2	11 ± 28
	Test	4.3 ± 1.9	55 ± 108	4.2 ± 3.2	10 ± 28
B	Train	**2.6 ± 1.2**	**25 ± 46**	**2.3 ± 2.1**	**5 ± 20**
	Val	**2.6 ± 1.2**	**25 ± 46**	**2.2 ± 2.1**	**5 ± 19**
	Test	**2.6 ± 1.2**	**25 ± 48**	**2.3 ± 2.1**	**5 ± 16**
C	Train	4.8 ± 2.4	**25 ± 45**	6.8 ± 4.3	**6 ± 19**
	Val	4.7 ± 2.4	**25 ± 44**	6.6 ± 4.3	**6 ± 19**
	Test	4.8 ± 2.4	**24 ± 44**	6.8 ± 4.3	**6 ± 19**

For geometry, three additional test cases with unseen cross-section shapes were considered: a square cuboid with $A = 5\,\mathrm{mm} \times 5\,\mathrm{mm}$, a rectangular cuboid with $A = 5\,\mathrm{mm} \times 10\,\mathrm{mm}$, and a cylinder with circular cross-section (diameter 5 mm). The corresponding error metrics are reported in Table 5, and representative thermal cycles are shown in Fig. 9. These tests quantify how far the trained model can be extrapolated in terms of process parameters and geometry while still yielding physically plausible temperature histories.

4. Results

4.1. Model-capacity sensitivity and selection

Influence of LSTM model capacity

Three LSTM architectures of different capacity (Models A–C) were evaluated to determine how the accuracy of the surrogate depends on model size and to select a representative configuration for the subsequent analyses. Model A comprises 1937 trainable parameters and consists of a single LSTM layer with 16 parallel cells. Trainable parameters denote the weights and biases of the network that are optimized during training and determine the model's representational capacity. Model B has 14 369 parameters distributed over two layers with 32 cells each, whereas Model C contains 119 873 parameters arranged in four layers with 64 cells per layer. Accordingly, the increase in parameter count from Model A to Model C reflects a systematic increase in model complexity and flexibility, enabling the networks to capture more intricate temporal dependencies in the thermal histories. For each architecture, an ensemble of five independently trained networks was generated. The training data used for model development is described in Section 3.2.1. From the ensemble predictions, mean values and confidence intervals were computed to quantify predictive uncertainty. Model quality was assessed using the root mean square error (RMSE) between predicted and FEM-simulated temperature–time histories, as well as process-specific error measures that capture relevant features of the thermal histories: the deviation of the maximum layer temperature $\Delta T_{\max}$, the deviation of the final temperature ΔT_{end}, and the deviation of the resulting hardness ΔHV. All error measures were computed for each evaluation point and subsequently aggregated over the complete dataset. Errors were evaluated separately for the training, validation, and test sets. Details for the FEM simulation and the deep learning setup are listed in Sections 3.1 and 3.2 respectively.

Table 2 summarizes the quantitative performance of the three architectures. Across all models, the errors for training, validation, and test sets are very similar, indicating good generalization and only minor overfitting.

Model A achieves an average RMSE of (4.4 ± 2.0) K and (4.3 ± 1.9) K for the training and test sets, respectively. The maximum temperature deviation $\Delta T_{\max}$ is on the order of (55 ± 105) K (train) and (55 ± 108) K (test), while the final-temperature error ΔT_{end} remains comparatively small at about (4.2 ± 3.2) K across all splits. The derived hardness deviation ΔHV is 10±28 HV1 on the test set and thus comparable to typical

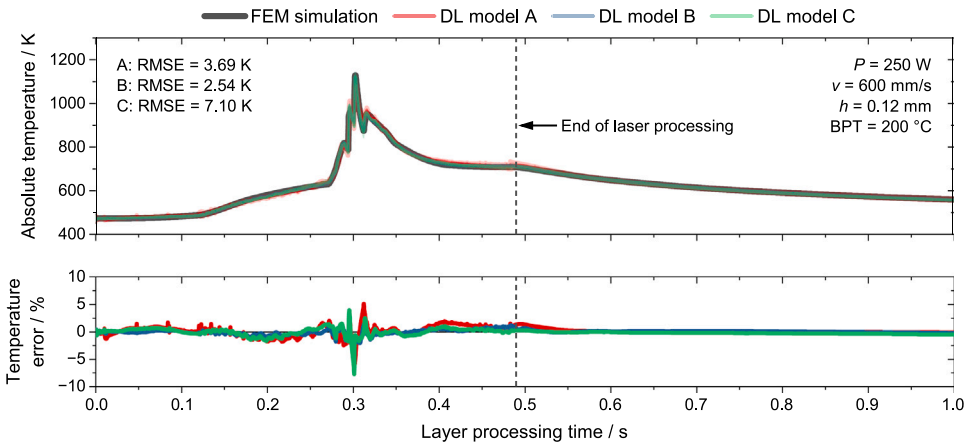

Fig. 4. Comparison of FEM-simulated and LSTM-predicted temperature histories for a representative near-surface evaluation point ($P = 250\,\mathrm{W}$, $v = 600\,\mathrm{mm\,s^{-1}}$, $h = 0.12\,\mathrm{mm}$, BPT $= 200\,°\mathrm{C}$, $0.2\,\mathrm{mm}$ beneath the current surface). Upper panel: absolute temperature as a function of layer processing time for the FEM reference and the three deep learning models A–C; ensemble RMSE values for each model are indicated. Lower panel: corresponding relative temperature error with respect to the FEM solution; the vertical dashed line denotes the end of laser processing.

experimental scatter. Taken together, these values already demonstrate a high level of consistency between the LSTM predictions and the FEM simulations, albeit with noticeable spread in the peak-temperature error.

With Model B, the prediction quality improves significantly and consistently for all datasets. The RMSE decreases to $(2.6\pm1.2)\,\mathrm{K}$ for training, validation, and test, corresponding to a reduction of roughly $40\,\%$ compared to Model A. At the same time, the maximum-temperature deviation $\Delta T_{\max}$ is nearly halved to about $(25\pm46)\,\mathrm{K}$ (train/val) and $(25\pm48)\,\mathrm{K}$ (test), indicating a much more accurate reconstruction of the thermal peaks. The final-temperature error ΔT_{end} drops to approximately $2.3\,\mathrm{K}$, i.e. an additional reduction of about $45\,\%$ relative to Model A. The hardness deviation ΔHV is reduced to 5 ± 16 HV1 on the test set, which represents an improvement of about $50\,\%$ and brings the mean error well within the experimental variability of hardness measurements. Importantly, both the mean values and standard deviations of all error metrics remain virtually unchanged from training to validation and test, underlining the robustness of Model B.

Further increasing the network size (Model C) does not yield additional performance gains. Although the maximum-temperature deviation $\Delta T_{\max}$ is comparable to that of Model B (around $(25\pm44)\,\mathrm{K}$ across all splits) and the hardness error ΔHV remains at a similar level (6 ± 19 HV1), the global fit of the temperature–time histories deteriorates. The RMSE rises to $(4.7\pm2.4)\,\mathrm{K}$ (validation) and $(4.8\pm2.4)\,\mathrm{K}$ (test), i.e. by about $85\,\%$ compared to Model B and even slightly above the error level of Model A. In addition, the final-temperature deviation ΔT_{end} increases considerably to $(6.6\pm4.3)\,\mathrm{K}$ (validation) and $(6.8\pm4.3)\,\mathrm{K}$ (test), indicating systematic inaccuracies in the late cooling regime. The larger mean errors together with increased standard deviations indicate that increasing the model capacity did not improve generalization under the applied training conditions and was associated with less stable convergence.

Within the three investigated LSTM capacity variants, Model B provides the most favorable balance of accuracy, stability and computational effort and was therefore selected as the reference model for the subsequent analyses.

Temperature and tempering parameter profiles

The quality of the predicted thermal histories and tempering response for a representative evaluation point is illustrated in Figs. 4 and 5. The example corresponds to a near-surface point $0.2\,\mathrm{mm}$ beneath the current surface processed with $P = 250\,\mathrm{W}$, $v = 600\,\mathrm{mm\,s^{-1}}$, $h = 0.12\,\mathrm{mm}$ and BPT $= 200\,°\mathrm{C}$.

Fig. 4 (top) shows the temperature–time profiles of the FEM reference solution and the three LSTM models. The thermal cycle is characterized by a rapid temperature increase during laser exposure, a pronounced peak when the laser passes the evaluation point, and a subsequent monotonic cooling phase. All three networks reproduce the timing and magnitude of these features with high fidelity. The predicted curves are almost indistinguishable from the FEM reference. The ensemble RMSE values of $3.89\,\mathrm{K}$, $2.54\,\mathrm{K}$, and $7.10\,\mathrm{K}$ for Models A, B, and C, respectively, quantify these differences. The lower panel shows the corresponding relative temperature error. For Model B, the error remains confined to a narrow band around zero over most of the layer processing time and exhibits only short, transient deviations in the vicinity of the temperature peak, where the local thermal gradients and cooling rates are highest. Models A and C display slightly larger deviations but no systematic bias over the full thermal cycle.

Fig. 5 presents the associated evolution of the Hollomon–Jaffe tempering parameter P_{HJ}. The FEM reference solution exhibits the expected three-stage behavior: a steep increase of P_{HJ} during heating (tempering of previously depositioned material), a quasi-stationary plateau while the material is in the austenitic or molten state, and a gradual decrease during cooling due to the martensitic transformation. The LSTM models accurately capture the onset and duration of each regime as well as the absolute level of P_{HJ}. In the lower panel, the relative P_{HJ} error remains small over the entire process time, with modest, short-lived deviations during the transition between the heating and plateau regimes and in the subsequent transformation-dominated cooling stage. Among the three candidates, Model B again yields the smallest deviations over the full tempering cycle. Taken together, these results demonstrate that the selected LSTM model reconstructs the transient temperature histories with high accuracy and reliably predicts the integrated tempering response that governs the local hardness.

Computational efficiency

The wall-clock times for training and inference of the three LSTM architectures are summarized in Table 3.For each architecture, the table lists the mean training time of a single network instance and the corresponding total time for an ensemble of five independently trained models, as well as the average inference times for batch and per-point evaluations (batch size 256). All runs were performed with identical training data, learning rate, and hyperparameters. Differences in runtime thus arise solely from the network complexity, i.e. the number of layers and cells.

As expected, the training time increases with model size. On average, a single instance of Model A requires $(8844\pm1767)\,\mathrm{s}$, whereas Model B requires $(9564\pm376)\,\mathrm{s}$, corresponding to a modest increase of about $8\,\%$. For the largest architecture, Model C, the training time

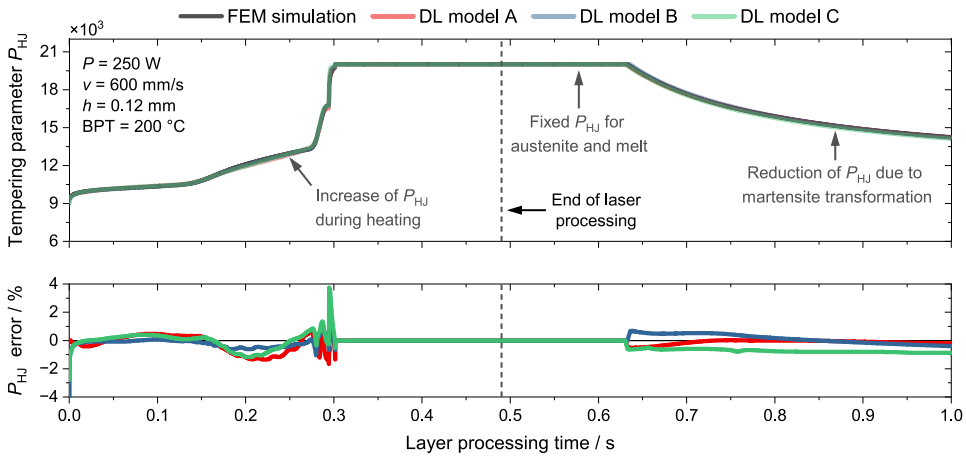

Fig. 5. Comparison of FEM-simulated and LSTM-predicted Hollomon–Jaffe tempering parameter P_{HJ} for the same evaluation point and process parameters as in Fig. 4. Upper panel: temporal evolution of P_{HJ}, illustrating the increase during heating, the quasi-stationary plateau in the austenitic/melt state, and the subsequent reduction during martensitic transformation. Lower panel: relative P_{HJ} error of the three deep learning models with respect to the FEM reference as a function of layer processing time.

Table 3
Training and inference times of the three LSTM models. For each architecture, the table reports the wall-clock training time of a single model instance and the total training time for an ensemble of five independently trained models, as well as the corresponding inference times for batched and per-point evaluations. All numerical entries are mean values with the associated standard deviation obtained from 5 repeated runs (given as mean ± std).

Model	Training time/s		Inference time/s	
	Single model	Ensemble	Ensemble[a]	Ensemble[b]
A	8844 ± 1767	44 220	11.4 ± 1.9	0.045 ± 0.007
B	9564 ± 376	47 820	11.5 ± 1.4	0.045 ± 0.006
C	11 916 ± 655	59 580	13.5 ± 1.3	0.053 ± 0.005

Note: All training and inference times were obtained with a batch size of 256.
[a] Average inference time per ensemble and batch.
[b] Average inference time per ensemble and evaluation point.

rises to (11 916 ± 655) s, i.e. approximately 25 % longer than Model B and 35 % longer than Model A. Because the ensemble members are trained sequentially in the present study, the total training time for the ensembles scales nearly linearly with these values and amounts to 44 220 s, 47 820 s, and 59 580 s for Models A, B, and C, respectively. This overhead could be reduced substantially by parallelizing the ensemble training on suitable hardware.

Differences in computational effort are also evident during inference, i.e. when applying the trained models to predict temperature histories. For a batch size of 256, the average ensemble inference times per batch are (11.4 ± 1.9) s for Model A, (11.5 ± 1.4) s for Model B, and (13.5 ± 1.3) s for Model C. These values correspond to average per-point inference times of (0.045 ± 0.007) s, (0.045 ± 0.006) s, and (0.053 ± 0.005) s, respectively. Within the uncertainty of the measurements, Models A and B exhibit essentially identical inference costs, whereas Model C is slower by about 18 % per evaluation point.

In contrast, the corresponding FEM simulations require on the order of 1×10^3 s to 1×10^5 s per equivalent temperature sequence. The LSTM surrogates thus provide speed-ups of several orders of magnitude, enabling rapid evaluation of large numbers of process parameter combinations and geometric positions. Taken together with the error metrics in Table 2, these results support the selection of the Model B ensemble as the reference configuration for all subsequent analyses within this study.

4.2. Analysis of the selected reference model

All analyses in this section are based on the ensemble of five networks of Model B. They are carried out exclusively on the test dataset, which was not used during training and therefore provides an independent assessment of the model performance.

4.2.1. Influence of input features

The influence of the individual input features on the prediction quality was quantified using the Permutation Feature Importance (PFI) method. The feature and data preparation is listed in Section 3.2.1. In this approach, the relative contribution of each feature is determined by randomly permuting its values while keeping all other inputs unchanged. The resulting change in RMSE of the temperature histories reflects the impact of the respective feature on the model accuracy. The input features comprise process parameters (laser power P, scan speed v, base plate temperature BPT, and laser on/off state), geometric descriptors of the relative position between laser and evaluation point (Δx_0, Δy_0, Δz_0, Δz_{BP}), temporal information ($\Delta t_{cooling}$, T_{t-1}), and the hatch angle represented by its cyclic components $\sin(\alpha_{SV})$ and $\cos(\alpha_{SV})$.

Fig. 6 illustrates the resulting hierarchy of feature importance. Out of the twelve input variables, ten lead to a substantial increase of the RMSE when perturbed, confirming their relevance for the prediction task. The strongest effects are observed for the relative distance between laser and evaluation point in y-direction Δy_0, the base plate temperature BPT, the laser power P, and the horizontal distance Δx_0. Permuting any of these variables increases the RMSE by roughly one order of magnitude compared to the baseline value, corresponding to a strong degradation of the temperature prediction. A similarly pronounced impact is found for the binary laser-on/off indicator, the time since the end of laser processing $\Delta t_{cooling}$, and the vertical distance between laser and evaluation point Δz_0, which underlines the importance of correctly representing the local process state and the thermal boundary conditions.

Variables related to the vertical positioning relative to the base plate and to the scan dynamics show an intermediate influence. Permutation of the distance between measurement point and base plate Δz_{BP} and of the scan speed v yields RMSE values clearly above the baseline, but below those associated with the primary process parameters P and BPT. These inputs mainly govern the temperature profile.

For the autoregressive temperature input T_{t-1}, the classical PFI procedure (random permutation) is not well posed, because T_{t-1} is deterministically coupled to the target temperature and naive shuffling would destroy the temporal structure of the sequence. Instead, T_{t-1} was

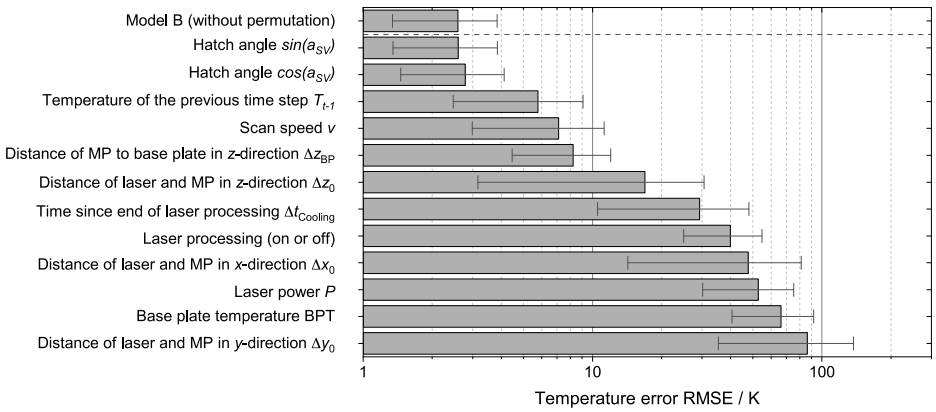

Fig. 6. Permutation Feature Importance (PFI) results for the Model B ensemble. The top bar shows the reference RMSE of Model B on the test dataset without permutation (baseline error of (2.6 ± 1.2) K). Each subsequent bar represents the RMSE obtained when the corresponding input variable is randomly permuted (or, for T_{t-1}, replaced by a constant value), while all other inputs are kept unchanged. Bars indicate mean values and whiskers the associated standard deviation over repeated permutations. The horizontal axis is plotted on a logarithmic scale.

fixed to a constant value and the resulting predictions were compared to the original model output. The corresponding bar in Fig. 6 shows an RMSE that is several times higher than the baseline, demonstrating the critical role of T_{t-1} for maintaining temporal consistency in the predicted temperature histories.

The individually permuted hatch-angle components $\sin(\alpha_{SV})$ and $\cos(\alpha_{SV})$ produce only a minor increase in RMSE, with values remaining close to the baseline within the indicated standard deviations. This result must be interpreted cautiously because both components jointly encode the same periodic variable and are therefore statistically dependent. Permuting only one component leaves part of the directional information available through the other component. Consequently, the individual PFI values may underestimate the combined relevance of the hatch-angle representation and do not, by themselves, justify removing these inputs.

Overall, the PFI analysis demonstrates that almost all process-related and geometric input variables contribute meaningfully to the performance of Model B, with a clear ranking from highly influential process parameters to moderately important geometric and kinematic quantities and, finally, to the hatch-angle components, which show low individual PFI scores within their coupled representation.

4.2.2. Influence of different process categories

The complete test dataset covers all geometry, scan strategy, and process variations and exhibits a global mean RMSE of (2.6 ± 1.2) K. All input variables considered in the PFI analysis were used for these evaluations. To investigate potential dependencies of the prediction accuracy on specific process conditions, this global error was decomposed into individual process categories. The analysis was carried out separately for process parameter combinations, base plate temperatures (BPT), scan vector lengths (equivalently, cross-sectional areas), and the distance of the evaluation points from the corresponding part surface (Fig. 7).

For the three investigated process parameter combinations, each comprising 20 000 data points, pronounced differences in model accuracy are observed (Fig. 7a). The parameter set P006 with the highest volumetric energy density (VED) exhibits the largest error with an RMSE of (3.2 ± 1.3) K. The corresponding boxplot is shifted to higher values and shows a wider interquartile range and a larger number of high-error outliers compared to the other parameter sets, indicating a broader spread of local prediction errors under strongly energy-intensive conditions. For P030 and P076, the RMSE decreases to (2.5 ± 1.1) K and (1.8 ± 0.7) K, respectively. The boxplots for these conditions are more compact, with lower medians and fewer extreme outliers, demonstrating that the model performs more accurately and more consistently under less extreme process conditions.

The dependence on base plate temperature is less pronounced (Fig. 7b). For the non-controlled BPT and the controlled BPT of 100 °C, the RMSE is (2.3 ± 1.1) K (each with 12 000 data points). The associated boxplots are almost identical in level and spread. At a controlled BPT of 200 °C, the RMSE slightly increases to (2.7 ± 1.2) K, and the distribution becomes marginally wider, suggesting a modest degradation of accuracy and a somewhat larger variability at higher preheating temperatures.

A clear trend is also observed with respect to scan vector length and thus cross-sectional area (each class containing 15 000 data points, Fig. 7c). For short scan vectors of 3 mm and 6 mm, the RMSE is (1.9 ± 1.0) K and (2.3 ± 0.9) K, respectively. The corresponding boxplots exhibit relatively narrow interquartile ranges and only a few larger outliers. For longer scan vectors of 9 mm and 12 mm, the errors increase to (3.1 ± 1.2) K and (2.7 ± 1.2) K, accompanied by broader boxes and more pronounced upper tails. This behavior indicates that more extended melt tracks with stronger geometric temperature gradients and longer interaction times are more challenging to predict accurately.

The distance of the evaluation points to the current part surface also has a marked influence on the prediction quality (Fig. 7d). Depending on the layer, between 50 and 5000 data points were evaluated. In near-surface regions, where temperatures above the liquidus temperature T_L occur and the thermal cycle is dominated by rapid heating and solidification, the boxplots are centered at higher RMSE values and display a pronounced right-hand tail. Here, the RMSE values are up to 25 % higher than the global mean. With increasing depth below the surface, the median and spread of the RMSE systematically decrease and reach values below 2.0 K at larger distances, reflecting the smoother and less transient thermal histories in the bulk material. Only a small number of outliers remains at deeper layers, indicating that large local deviations are predominantly confined to the highly transient near-surface region.

4.3. Generalization of the selected reference model

To assess the generalization capability of the trained deep learning model, selected input variables were deliberately varied beyond the ranges used during training. The analysis covers both process parameters and the geometry of the manufactured cross sections.

Variation of process parameters

During training, five discrete hatch angles were used for the cyclically transformed variables $\sin(\alpha_{SV})$ and $\cos(\alpha_{SV})$ (0°, 18°, 36°, 54° and 72°). Although the hatch-angle components showed a low importance in the PFI analysis, this result is affected by the coupled sine–cosine representation (see Section 4.2.1). The hatch angle was therefore varied

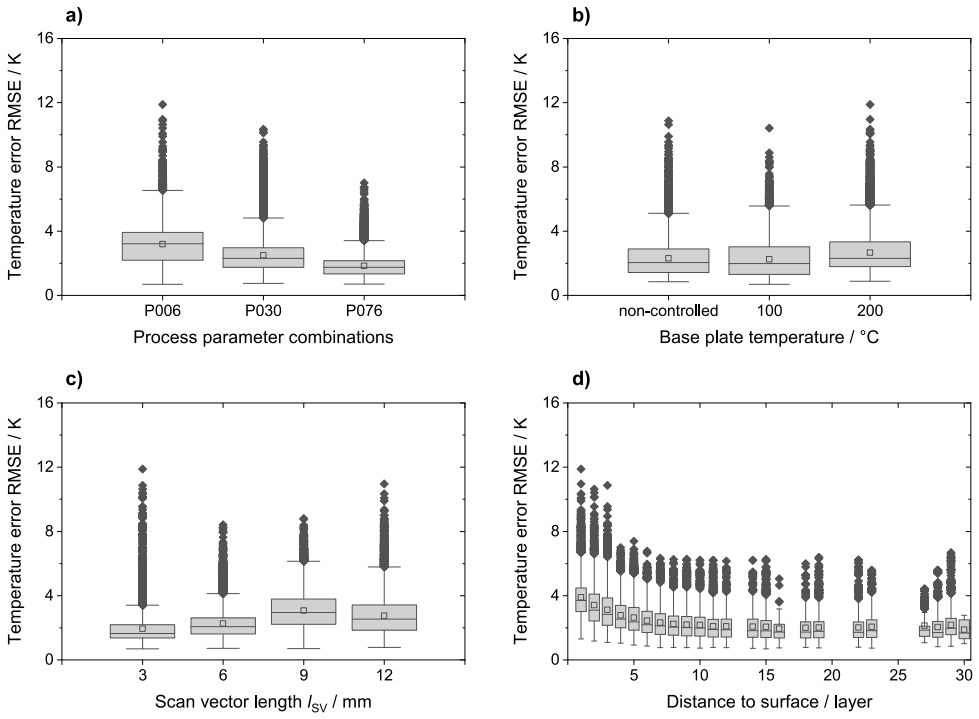

Fig. 7. RMSE of the Model B ensemble as a function of process categories: (a) process parameter combination, (b) base plate temperature BPT, (c) scan vector length l_{SV} and (d) distance to the current part surface. Boxplots indicate the median and interquartile range, whiskers denote the 1.5× interquartile range, and symbols represent individual evaluation points outside the interquartile range. Each point corresponds to the RMSE of the temperature history at a single evaluation point.

Table 4
Error metrics for variations of process parameters compared with the reference dataset. All values are given as mean ± standard deviation.

Dataset	Error metrics			
	RMSE/K	ΔT_{max}/K	ΔT_{end}/K	ΔHV/HV1
Reference	2.2 ± 0.8	25 ± 40	0.9 ± 0.5	3 ± 21
Variation α_{SV}	2.2 ± 1.0	25 ± 41	1.8 ± 0.8	3 ± 12
Variation BPT	2.3 ± 0.9	23 ± 38	1.7 ± 1.0	3 ± 19
Variation P	8.6 ± 5.4	119 ± 221	1.6 ± 0.8	20 ± 42
Variation v	11.9 ± 6.4	140 ± 201	2.3 ± 1.3	23 ± 45
Variation α_{SV}, BPT	2.0 ± 0.8	20 ± 33	1.8 ± 0.7	3 ± 20
Variation α_{SV}, P	8.5 ± 4.8	120 ± 226	1.6 ± 0.8	20 ± 38
Variation α_{SV}, v	12.1 ± 6.5	140 ± 201	2.3 ± 1.3	25 ± 48

explicitly to assess whether directional effects are nevertheless captured by the model. In addition, three discrete base plate temperatures (BPTs) and three process parameter combinations of laser power, scan speed, and hatch spacing were included in the training data. For the generalization study, the input variables in the test dataset were sampled randomly such that both interpolation within and extrapolation beyond the training ranges could be evaluated.

The resulting error metrics are summarized in Table 4. For randomly selected hatch angles (17°, 41°, 46°, 56°, 75° and 89°) that were not part of the training dataset, the errors are virtually identical to the reference case. The RMSE remains at (2.2 ± 1.0) K, matching the level of the reference dataset with (2.2 ± 0.8) K. The deviation of the maximum temperature ΔT_{max} is (25 ± 41) K, the deviation of the final temperature ΔT_{end} is (1.8 ± 0.8) K, and the hardness difference ΔHV is (3 ± 12) HV1. A similar behavior is observed for variations of the BPT and the combined variation of hatch angle and BPT, with RMSE values of (2.3 ± 0.9) K and (2.0 ± 0.8) K, respectively. These results indicate that the model interpolates reliably within and between the discrete process states seen during training.

In contrast, variations of laser power and scan speed, alone or in combination with the hatch angle, lead to significantly higher errors. For these extrapolative changes, the RMSE increases to 8.5 K to 12.1 K on average, and the deviation of the maximum temperature ΔT_{max} rises from 25 K (reference) to values between 119 K to 140 K. The hardness difference ΔHV grows from about 3 HV1 to values between 20 HV1 to 25 HV1, while the deviation of the final temperature ΔT_{end} changes only moderately, from 0.9 K to 1.6 K to 2.3 K. These trends indicate that the model is more sensitive to strong changes in the absorbed energy density and the local solidification conditions than to variations of purely geometric or boundary-condition-type inputs.

Fig. 8 illustrates representative temperature–time histories for two such extrapolative cases, obtained by varying the scan speed. For this purpose, modified process parameter combinations that were not contained in the training data were evaluated using the deep learning model (inference). Corresponding FEM simulations were performed for validation.

For a reduced scan speed of 360 mm s^{-1} (Fig. 8a) and an increased scan speed of 440 mm s^{-1} (Fig. 8b), the deep learning model and FEM simulations exhibit very similar thermal cycles. The temporal position and magnitude of the main temperature peak and the subsequent cooling branch are reproduced with good accuracy. The local RMSE for the representative points shown is on the order of 6 K to 7 K, in agreement with the ensemble statistics in Table 4. Despite the increased RMSE for strong variations of v, the predicted temperature histories remain physically consistent and yield hardness deviations that are still acceptable for process-screening applications.

Variation of cross-section geometry

In addition to process parameters, the generalization capability was also investigated by targeted variation of the cross-section geometry. All training data were based on cuboid specimens with square cross-sections A, such that a simple relationship between cross-sectional area and the square of the scan vector length could be assumed ($A = l_{SV}^2$). To test the model beyond this setting, three new test datasets with

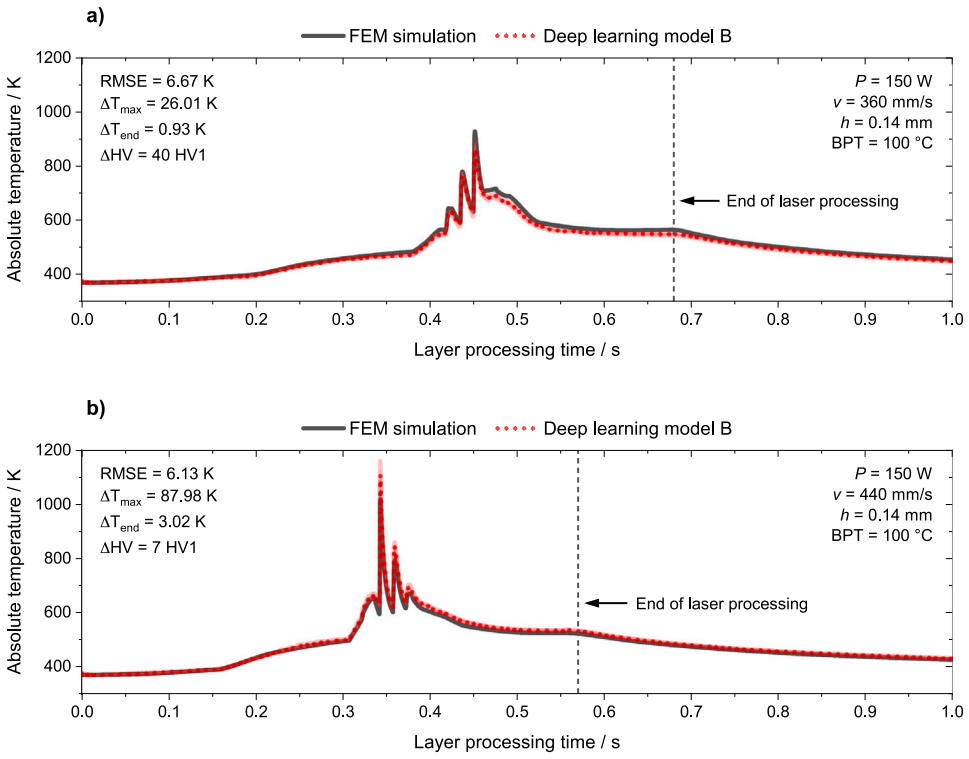

Fig. 8. Temperature–time histories of representative evaluation points for variations of the scan speed used to assess the generalization capability. Direct comparison between the predictions of the deep learning model (Model B) and the FEM simulations for (a) a reduced scan speed of 360 mm s^{-1} and (b) an increased scan speed of 440 mm s^{-1}. The reported RMSE and hardness deviations ΔHV correspond to the specific evaluation points shown. The modified process parameter combinations were not part of the training, validation, or test datasets.

Table 5
Error metrics of the selected reference model for variations of the cross-section area A. The reference dataset corresponds to the original cuboid specimens with square cross-section used for training. Geometry variation I denotes a cuboid with square cross-section $A = 5\,\text{mm} \times 5\,\text{mm}$ (interpolation), variation II a cuboid with rectangular cross-section $A = 5\,\text{mm} \times 10\,\text{mm}$ (extrapolation in aspect ratio), and variation III a cylinder with circular cross-section of 5 mm diameter (extrapolation in shape). All error metrics are given as mean ± standard deviation.

Dataset	Error metrics			
	RMSE/K	ΔT_{max}/K	ΔT_{end}/K	ΔHV/HV1
Reference	2.2 ± 0.8	25 ± 40	0.9 ± 0.5	3 ± 21
Geometry variation I	3.3 ± 0.9	27 ± 45	3.4 ± 1.2	5 ± 21
Geometry variation II	5.9 ± 2.5	50 ± 87	1.8 ± 1.4	9 ± 26
Geometry variation III	4.0 ± 1.1	34 ± 56	3.8 ± 1.4	6 ± 27

modified cross-sectional geometries were defined: (i) a cuboid with square cross-section $A = 5\,\text{mm} \times 5\,\text{mm}$, representing interpolation within the range of l_{SV} (geometry variation I); (ii) a cuboid with rectangular cross-section $A = 5\,\text{mm} \times 10\,\text{mm}$, corresponding to extrapolation in cross-section aspect ratio at approximately fixed area range (geometry variation II); and (iii) a cylinder with circular cross-section and a diameter of 5 mm, representing extrapolation in both shape and area distribution (geometry variation III).

The results in Table 5 show that changes in geometry lead to moderate, but systematic, variations in the error metrics. For geometry variation I, the RMSE increases from (2.2 ± 0.8) K (reference) to (3.3 ± 0.9) K. The deviations of the maximum temperature ΔT_{max} and the resulting hardness ΔHV remain within the standard deviation of the reference dataset, while the final-temperature error ΔT_{end} increases to (3.4 ± 1.2) K, compared to (0.9 ± 0.5) K for the reference. This indicates that the model interpolates reasonably well to intermediate square cross-sections, but exhibits a slightly larger bias in the late cooling regime.

For geometry variation II (rectangular cross-section), the RMSE rises to (5.9 ± 2.5) K. The deviation of the maximum temperature increases to (50 ± 87) K, and the mean hardness deviation ΔHV reaches 9 HV1. Geometry variation III (cylindrical geometry) lies between the other two variants, with an RMSE of (4.0 ± 1.1) K, a maximum-temperature deviation ΔT_{max} of (34 ± 56) K, and a hardness deviation ΔHV of 6 HV1. The final-temperature error ΔT_{end} amounts to (3.8 ± 1.4) K in this case. Overall, the largest errors occur for the most pronounced extrapolation in terms of aspect ratio (geometry variation II), while the purely shape-related extrapolation (cylinder) leads to intermediate deviations.

Fig. 9 illustrates representative temperature–time histories for geometry variations I–III. For all tested geometries, the temperature–time histories predicted by the deep learning model and the FEM simulations agree well both qualitatively and quantitatively: the timing and magnitude of the principal temperature peak and the overall cooling trend are reproduced with good fidelity. Small discrepancies occur primarily in the near field of the laser–material interaction and during the late cooling phase, but no physically inconsistent behavior (e.g. spurious oscillations or non-monotonic cooling) is observed. These results demonstrate that the selected LSTM model retains a reasonable level of accuracy when applied to previously unseen cross-section geometries, particularly for applications where moderate errors in peak temperature and hardness are acceptable.

5. Discussion

The results of the deep learning based process modeling demonstrate that recurrent neural networks, in particular LSTM architectures, can predict the time-resolved temperature evolution in PBF-LB with high accuracy and good physical consistency. In the following, the main findings are discussed in the context of the numerical simulations, experimental observations and current literature.

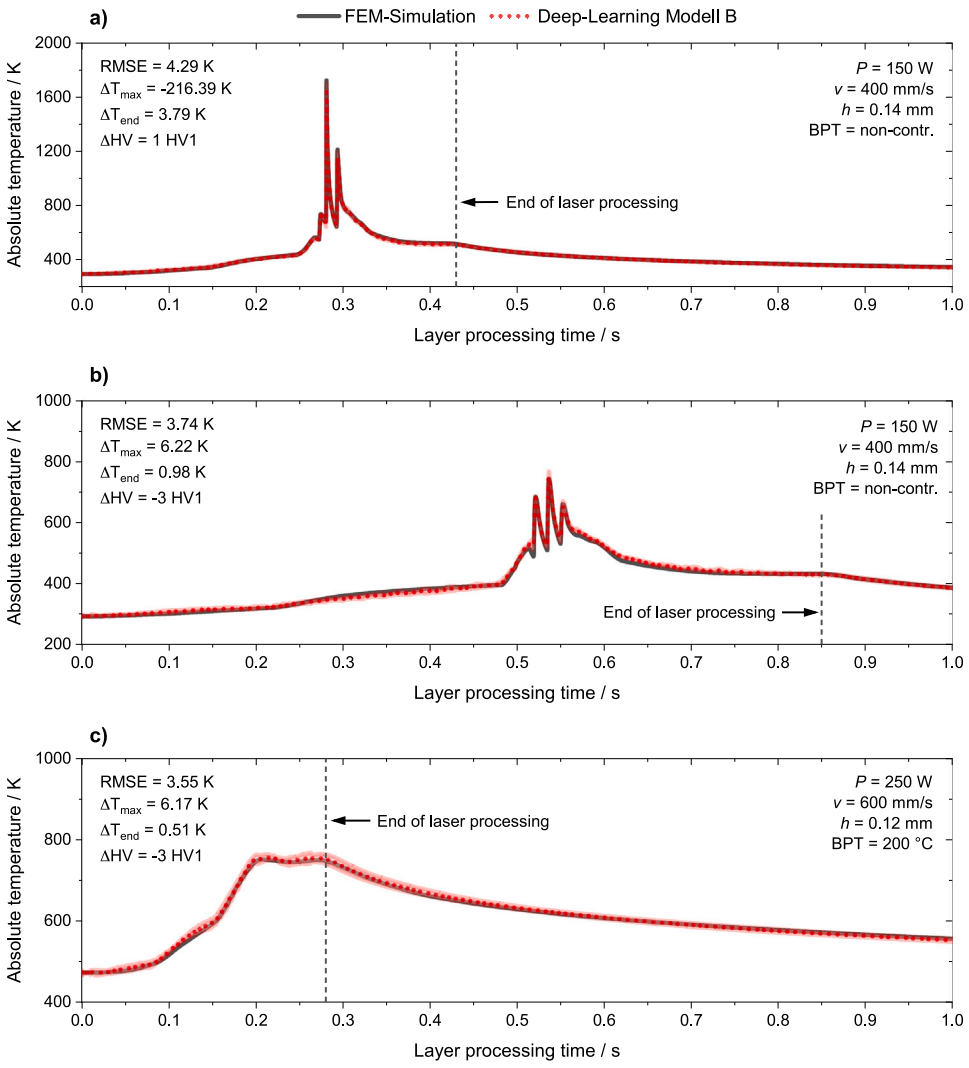

Fig. 9. Temperature–time histories of a representative evaluation point for the three cross-section geometries considered in the generalization study: (a) square cuboid with $A = 5\,\text{mm} \times 5\,\text{mm}$ (geometry variation I), (b) rectangular cuboid with $A = 5\,\text{mm} \times 10\,\text{mm}$ (geometry variation II), and (c) cylinder with circular cross-section, diameter 5 mm (geometry variation III). In all cases, the same process parameter combinations as in the reference dataset are used. Solid lines denote the FEM simulations and dotted lines the predictions of the deep learning model (Model B). Each curve pair illustrates the local agreement over the full layer processing time, including rapid heating during laser exposure, the main temperature peak, and the subsequent cooling phase. The inset values (RMSE, ΔT_{max}, ΔT_{end}, ΔHV) refer to the specific evaluation points shown and are consistent with the ensemble statistics reported in Table 5.

Comparison with FEM simulations and experimental observations

The predictions of the selected LSTM ensemble (Model B) show very good agreement with the FEM simulations in terms of shape, amplitude and temporal dynamics of the thermal cycles (Figs. 4 and 5). Both the rapid heating during laser exposure and the subsequent cooling phase are reproduced in a physically plausible manner. The mean deviations of the temperature histories are small, with RMSE values around (2.6 ± 1.2) K (Table 2). The hardness values derived via the Hollomon–Jaffe tempering parameter are likewise in very good agreement with the FEM-based reference and remain within the typical experimental scatter of hardness measurements. For the selected reference model, the average hardness error over the full test set is below 20 HV1.

Larger discrepancies are mainly observed due to deviations in the near-field region of the laser track, where small temporal shifts in the maximum temperature lead to disproportionately large differences in phase fractions and hence hardness (Fig. 8). Nevertheless, the model consistently reproduces the experimentally observed trends between volumetric energy density, base plate temperature and resulting hardness, capturing the reduction in hardness scatter at higher energy input and preheating. Taken together, these results confirm that the LSTM ensemble provides a reliable surrogate for computationally expensive thermal FEM simulations while retaining sufficient fidelity to be used for microstructure- and property-related analyses.

Assessment of model capacity and input features

Within the investigated LSTM capacity variants, the intermediate configuration (Model B) offers the most favorable balance of prediction accuracy, numerical stability and computational cost. Increasing the number of layers and cells in Model C did not improve the error metrics under the applied training conditions and was associated with less stable convergence, as reflected by the higher RMSE, the larger spread of the final-temperature error and the increased training time.

The PFI analysis (Fig. 6) reveals that almost all input variables contribute significantly to the model performance. Laser power, base plate temperature and the relative position between laser and evaluation point (Δx_0, Δy_0, Δz_0) are the most influential features, which is consistent with their direct control over local energy input and heat conduction. The hatch-angle components show low individual PFI

scores. However, their coupled sine–cosine representation means that these scores may underestimate their combined relevance. The very strong sensitivity with respect to the autoregressive temperature input T_{t-1} confirms the importance of temporal correlations for learning the thermal dynamics: fixing T_{t-1} to a constant value leads to a pronounced increase in RMSE. Overall, these findings support the chosen physics-inspired feature design, in which distances, process state and energy input are explicitly represented and exploited by the LSTM model.

Generalization and transferability

The generalization analyses indicate that the selected reference model generalizes well to interpolated and selected extrapolated process conditions. For variations of hatch angle and base plate temperature within and slightly beyond the training range, the global error metrics remain close to those of the reference dataset (Table 4). Similarly, changes in cross-section geometry lead to moderate increases in RMSE and hardness error, even for a cylindrical specimen that was not present in the training data, the RMSE remains on the order of 4 K (Table 5, Fig. 9). In all these cases, the predicted temperature–time histories remain physically consistent and reproduce the characteristic heating, peak and cooling behavior.

Higher errors occur when laser power or scan speed are varied beyond the three discrete parameter combinations used for training. Here, the RMSE increases to 8.5 K to 12.1 K and the error at maximum temperature exceeds 100 K on average (Table 4), although the qualitative shape of the thermal cycle is still captured very well (Fig. 8). This behavior is consistent with the dominant role of energy input in governing the thermal response and indicates that the model correctly captures the underlying physical dependencies, while remaining sensitive to variations that are underrepresented in the training data. The training dataset was intentionally limited to a small number of discrete process parameter combinations in order to assess the feasibility of the approach under constrained data conditions.

The results demonstrate that, even within this restricted setting, the model provides robust and physically consistent predictions. A systematic extension of the training space to include additional combinations of laser power, scan speed and hatch spacing is therefore expected to further improve quantitative accuracy, without changing the overall modeling approach.

Relation to existing approaches

Previously published thermal surrogates address complementary formulations. GRU- and RNN–DNN-based models represent pointwise histories or scan-pattern-dependent thermal fields [17,18], physics-engineered tree models seek transferability without an explicitly recurrent architecture [19], TCNs use causal temporal convolutions [22], and RGNNs emphasize mesh connectivity and transfer across geometries and materials [20,21]. Classical PINNs instead incorporate governing equations and boundary or initial conditions into the training objective, with demonstrations ranging from single-track PBF-LB to multilayer DED [24–26]. The recent STEAM framework further demonstrates scalable part-level scanwise thermal prediction for PBF-LB using spatial and temporal physics-based features [23]. Among published approaches, STEAM provides the closest quantitative PBF-LB comparison because it evaluated LSTM and GRU surrogates against FE-generated scanwise thermal histories on a common dataset [23]. On its large-scale unseen test geometries, the best-performing medium GRU achieved an average RMSE of 49.35 K, whereas the investigated LSTM variants yielded 82.77–124.71 K. The present LSTM ensemble yields 2.6 K on the independent test set and 3.3–5.9 K for unseen cross-section geometries. The temporal discretizations also differ: STEAM uses nonuniform increments with 0.1 ms during heating and temperature output every 180 ms during cooling, whereas the present model uses a fixed increment of 0.2 ms. The substantially lower errors in the present study therefore provide favorable literature context, but not a controlled ranking: STEAM predicts millions of nodes in geometries up to approximately ten times larger than its training cases, and the studies differ in material, thermal model, spatial representation and evaluation domain. The comparison nevertheless shows that the compact recurrent model provides competitive accuracy for the local process–property prediction task addressed here.

For the present process–property prediction task, the LSTM ensemble combines the accuracy discussed above with a substantial reduction in computational cost: inference at a single evaluation point requires only a few tens of milliseconds, compared to 10×10^3 s to 10×10^5 s for the corresponding FEM simulation. By incorporating physics-informed input features (energy input, relative laser position, distance to the free surface, autoregressive temperature), the model implicitly encodes the dominant process physics without explicitly solving the governing equations. The prediction task also differs from applications focused on melt-track geometry or density [27] and from direct data-driven prediction of hardness or mechanical properties [30,31]. Earlier hybrid process–property studies combined simulated or measured thermal information with mechanistic relations to estimate microstructure and microhardness, but were demonstrated for deposited Ni-based systems rather than PBF-LB of a martensitic quenched-and-tempered steel [28,29]. Here, the network first reconstructs the complete local temperature history. The resulting phase-transformation and non-isothermal tempering state and hardness are then calculated through the validated Hollomon–Jaffe relationship. The model does not predict geometry itself, but its transferability is evaluated under previously unseen cross-section geometries. This thermal-history intermediate preserves information on heating and cooling rates, peak temperatures and repeated exposure and allows the surrogate to remain coupled to existing transformation and property models.

Overall, the present study demonstrates a computationally efficient process–structure–property framework that connects process and geometric inputs to complete local thermal histories and derived hardness, with promising perspectives for accelerated process optimization and, in the longer term, integration into real-time process control.

6. Conclusion

A physics-informed deep learning surrogate for PBF-LB of AISI 4140 was developed to reconstruct complete local thermal histories and propagate them through a non-isothermal tempering model to predict hardness. The framework was validated against high-fidelity multiscale FE simulations. The sequence-to-sequence LSTM ensemble reconstructs local temperature–time histories with a mean RMSE of (2.6 ± 1.2) K, while the deviation of maximum temperature is (25 ± 48) K and the final-temperature error is (2.3 ± 2.1) K. The resulting hardness predictions show deviations of 5 ± 16 HV1, remaining within the typical experimental scatter for hardness measurements.

Within the investigated LSTM capacity variants, the two-layer Model B with 14 369 parameters provides stable performance across training, validation, and test sets and was selected as the reference model. Inference requires approximately 0.045 s per evaluation point, compared to 1×10^3 s to 1×10^5 s for equivalent FE simulations, corresponding to a speed-up of 4 to 6 orders of magnitude.

The feature importance analysis shows that permuting key inputs such as laser power, base plate temperature, or spatial laser–point distances increases the RMSE by up to one order of magnitude relative to the baseline of 2.6 K. Fixing the autoregressive input T_{t-1} similarly leads to a multi-fold increase in error, confirming its central role in capturing temporal evolution.

For interpolative generalization, variations in hatch angle and base plate temperature yield RMSE values of 2.0 K to 2.3 K, with maximum temperature deviations of approximately 20 K to 41 K and hardness deviations of 3 HV1 to 20 HV1. For extrapolative variations in laser power and scan speed, the RMSE increases to 8.5 K to 12.1 K, while

ΔT_{max} reaches up to 140 K and hardness deviations increase to 20 HV1 to 25 HV1. For unseen geometries, RMSE values range from 3.3 K (interpolated square cross-section) to 5.9 K (rectangular cross-section), with corresponding hardness deviations between 5 HV1 and 9 HV1.

Overall, the presented framework connects process and geometric inputs to complete local thermal histories and derived non-isothermal tempering hardness with millisecond-level inference times, providing a data-efficient pathway toward real-time-capable process modeling and digital twins in metal additive manufacturing.

CRediT authorship contribution statement

Philipp Schüßler: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Volker Schulze:** Writing – review & editing, Supervision. **Stefan Dietrich:** Writing – review & editing, Supervision, Project administration, Methodology, Conceptualization.

Declaration of competing interest

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

Acknowledgments

This study was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation, grant number 516837935). The authors also acknowledge the support by the KITPublication Fund of the Karlsruhe Institute of Technology (KIT).

Data availability

The FEM-based datasets used for training, validation, and testing are available in the KITopen repository of the Karlsruhe Institute of Technology (KIT) [46]. The trained LSTM ensemble models [47] as well as the corresponding code repository [48] are also publicly accessible via KITopen.

References

[1] S.F. Nabavi, H. Dalir, A. Farshidianfar, A comprehensive review of recent advances in laser powder bed fusion characteristics modeling: metallurgical and defects, Int. J. Adv. Manuf. Technol. 132 (5) (2024) 2233–2269, http://dx.doi.org/10.1007/s00170-024-13491-1.

[2] W. Hearn, S. Goel, C. Pauzon, A.S. Shaikh, N. Casati, E. Hryha, S. Van Petegem, Phase transformations in hypoeutectoid carbon steel during laser powder bed fusion, Addit. Manuf. 113 (2025) 105020, http://dx.doi.org/10.1016/j.addma.2025.105020.

[3] M. Markl, C. Körner, Multiscale modeling of powder bed–based additive manufacturing, Annu. Rev. Mater. Res. 46 (Volume 46, 2016) (2016) 93–123, http://dx.doi.org/10.1146/annurev-matsci-070115-032158.

[4] M.F. Zaeh, G. Branner, Investigations on residual stresses and deformations in selective laser melting, Prod. Eng. 4 (1) (2010) 35–45, http://dx.doi.org/10.1007/s11740-009-0192-y.

[5] E.R. Denlinger, J. Irwin, P. Michaleris, Thermomechanical modeling of additive manufacturing large parts, J. Manuf. Sci. Eng. 136 (6) (2014) http://dx.doi.org/10.1115/1.4028669.

[6] S. Hocine, H. van Swygenhoven, S. van Petegem, Verification of selective laser melting heat source models with operando X-ray diffraction data, Addit. Manuf. 37 (2021) 101747, http://dx.doi.org/10.1016/j.addma.2020.101747.

[7] Z. Luo, Y. Zhao, A survey of finite element analysis of temperature and thermal stress fields in powder bed fusion additive manufacturing, Addit. Manuf. 21 (2018) 318–332, http://dx.doi.org/10.1016/j.addma.2018.03.022.

[8] D. Sarkar, A. Kapil, A. Sharma, Advances in computational modeling for laser powder bed fusion additive manufacturing: A comprehensive review of finite element techniques and strategies, Addit. Manuf. 85 (2024) 104157, http://dx.doi.org/10.1016/j.addma.2024.104157.

[9] P. Schüßler, J. Damon, F. Mühl, S. Dietrich, V. Schulze, Laser surface hardening: A simulative study of tempering mechanisms on hardness and residual stress, Comput. Mater. Sci. 221 (2023) 112079, http://dx.doi.org/10.1016/j.commatsci.2023.112079.

[10] J. Liu, J. Ye, D. Silva Izquierdo, A. Vinel, N. Shamsaei, S. Shao, A review of machine learning techniques for process and performance optimization in laser beam powder bed fusion additive manufacturing, J. Intell. Manuf. 34 (8) (2023) 3249–3275, http://dx.doi.org/10.1007/s10845-022-02012-0.

[11] M.D. Xames, F.K. Torsha, F. Sarwar, A systematic literature review on recent trends of machine learning applications in additive manufacturing, J. Intell. Manuf. 34 (6) (2023) 2529–2555, http://dx.doi.org/10.1007/s10845-022-01957-6.

[12] K. Chen, P. Zhang, H. Yan, G. Chen, T. Sun, Q. Lu, Y. Chen, H. Shi, A review of machine learning in additive manufacturing: design and process, Int. J. Adv. Manuf. Technol. 135 (3) (2024) 1051–1087, http://dx.doi.org/10.1007/s00170-024-14543-2.

[13] M. Faegh, S. Ghungrad, J.P. Oliveira, P. Rao, A. Haghighi, A review on physics-informed machine learning for process-structure-property modeling in additive manufacturing, J. Manuf. Process. 133 (2025) 524–555, http://dx.doi.org/10.1016/j.jmapro.2024.11.066.

[14] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735–1780, http://dx.doi.org/10.1162/neco.1997.9.8.1735.

[15] K. Cho, B. van Merriënboer, D. Bahdanau, Y. Bengio, On the properties of neural machine translation: Encoder–decoder approaches, in: D. Wu, M. Carpuat, X. Carreras, E.M. Vecchi (Eds.), Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Association for Computational Linguistics, Doha, Qatar, 2014, pp. 103–111, http://dx.doi.org/10.3115/v1/W14-4012.

[16] J. Chung, C. Gulcehre, K. Cho, Y. Bengio, Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014, http://dx.doi.org/10.48550/arXiv.1412.3555, arXiv:1412.3555.

[17] M. Mozaffar, A. Paul, R. Al-Bahrani, S. Wolff, A. Choudhary, A. Agrawal, K. Ehmann, J. Cao, Data-driven prediction of the high-dimensional thermal history in directed energy deposition processes via recurrent neural networks, Manuf. Lett. 18 (2018) 35–39, http://dx.doi.org/10.1016/j.mfglet.2018.10.002.

[18] K. Ren, Y. Chew, Y.F. Zhang, J.Y.H. Fuh, G.J. Bi, Thermal field prediction for laser scanning paths in laser aided additive manufacturing by physics-based machine learning, Comput. Methods Appl. Mech. Engrg. 362 (2020) 112734, http://dx.doi.org/10.1016/j.cma.2019.112734.

[19] K.L. Ness, A. Paul, L. Sun, Z. Zhang, Towards a generic physics-based machine learning model for geometry invariant thermal history prediction in additive manufacturing, J. Mater. Process. Technol. 302 (2022) 117472, http://dx.doi.org/10.1016/j.jmatprotec.2021.117472.

[20] M. Mozaffar, S. Liao, H. Lin, K. Ehmann, J. Cao, Geometry-agnostic data-driven thermal modeling of additive manufacturing processes using graph neural networks, Addit. Manuf. 48 (2021) 102449, http://dx.doi.org/10.1016/j.addma.2021.102449.

[21] J.Y. Choi, S.M. Estalaki, D. Quispe, R. Zha, R. Rolark, M. Mozaffar, J. Cao, Transfer learning enabled geometry, process, and material agnostic RGNN for temperature prediction in directed energy deposition, Addit. Manuf. 109 (2025) 104876, http://dx.doi.org/10.1016/j.addma.2025.104876.

[22] V. Perumal, D. Abueidda, S. Koric, A. Kontsos, Temporal convolutional networks for data-driven thermal modeling of directed energy deposition, J. Manuf. Process. 85 (2023) 405–416, http://dx.doi.org/10.1016/j.jmapro.2022.11.063.

[23] B. Bostan, P. Vulimiri, S. Hinnebusch, D. Aryal, A.C. To, STEAM: A scalable data-driven surrogate modeling framework for part-scale scanwise thermal process simulation of laser powder bed fusion, Addit. Manuf. 118 (2026) 105090, http://dx.doi.org/10.1016/j.addma.2026.105090.

[24] Q. Zhu, Z. Liu, J. Yan, Machine learning for metal additive manufacturing: predicting temperature and melt pool fluid dynamics using physics-informed neural networks, Comput. Mech. 67 (2) (2021) 619–635, http://dx.doi.org/10.1007/s00466-020-01952-9.

[25] E. Hosseini, P. Scheel, O. Müller, R. Molinaro, S. Mishra, Single-track thermal analysis of laser powder bed fusion process: Parametric solution through physics-informed neural networks, Comput. Methods Appl. Mech. Engrg. 410 (2023) 116019, http://dx.doi.org/10.1016/j.cma.2023.116019.

[26] B. Peng, A. Panesar, Multi-layer thermal simulation using physics-informed neural network, Addit. Manuf. 95 (2024) 104498, http://dx.doi.org/10.1016/j.addma.2024.104498.

[27] M. Kuehne, K. Bartsch, B. Bossen, C. Emmelmann, Predicting melt track geometry and part density in laser powder bed fusion of metals using machine learning, Prog. Addit. Manuf. 8 (1) (2023) 47–54, http://dx.doi.org/10.1007/s40964-022-00387-3.

[28] Z. Gan, H. Li, S.J. Wolff, J.L. Bennett, G. Hyatt, G.J. Wagner, J. Cao, W.K. Liu, Data-driven microstructure and microhardness design in additive manufacturing using a self-organizing map, Engineering 5 (4) (2019) 730–735, http://dx.doi.org/10.1016/j.eng.2019.03.014.

[29] S.J. Wolff, Z. Gan, S. Lin, J.L. Bennett, W. Yan, G. Hyatt, K.F. Ehmann, G.J. Wagner, W.K. Liu, J. Cao, Experimentally validated predictions of thermal history and microhardness in laser-deposited Inconel 718 on carbon steel, Addit. Manuf. 27 (2019) 540–551, http://dx.doi.org/10.1016/j.addma.2019.03.019.

[30] X. Xie, J. Bennett, S. Saha, Y. Lu, J. Cao, W.K. Liu, Z. Gan, Mechanistic data-driven prediction of as-built mechanical properties in metal additive manufacturing, Npj Comput. Mater. 7 (1) (2021) 86, http://dx.doi.org/10.1038/s41524-021-00555-z.

[31] A. Yu, Y. Pan, L. Yue, F. Kang, J. Zhang, X. Lu, Explainable machine learning for hardness prediction of laser powder bed fused Ti-6Al-4V and assisting in understanding effects of process parameters, J. Alloys Compd. 1008 (2024) 176566, http://dx.doi.org/10.1016/j.jallcom.2024.176566.

[32] J.H. Hollomon, L.D. Jaffe, Time-temperature relations in tempering steel, Trans. AIME 162 (1945) 223–249.

[33] C. Gomes, A.-L. Kaiser, J.-P. Bas, A. Aissaoui, M. Piette, Predicting the mechanical properties of a quenched and tempered steel thanks to a ‘‘tempering parameter’’, Rev. Métall. 107 (7–8) (2010) 293–302, http://dx.doi.org/10.1051/metal/2010061.

[34] P. Schüßler, V. Schulze, S. Dietrich, Deep learning based prediction of thermal history during the laser-based powder bed fusion (PBF-LB) additive manufacturing process, in: Solid Freeform Fabrication 2024, University of Texas at Austin, Austin, TX, 2024, http://dx.doi.org/10.26153/tsw/58117.

[35] P. Schüßler, N. Nouri, V. Schulze, S. Dietrich, A novel multiscale process simulation to predict the impact of intrinsic heat treatment on local microstructure gradients and bulk hardness of AISI 4140 manufactured by laser powder bed fusion, Virtual Phys. Prototyp. 18 (1) (2023) e2271455, http://dx.doi.org/10.1080/17452759.2023.2271455.

[36] Q. Zhang, J. Xie, Z. Gao, T. London, D. Griffiths, V. Oancea, A metallurgical phase transformation framework applied to SLM additive manufacturing processes, Mater. Des. 166 (2019) 107618, http://dx.doi.org/10.1016/j.matdes.2019.107618.

[37] T. Mioković, J. Schwarzer, V. Schulze, O. Vöhringer, D. Löhe, Description of short time phase transformations during the heating of steels based on high-rate experimental data, J. Phys. IV 120 (2004) 591–598, http://dx.doi.org/10.1051/jp4:2004120068.

[38] S.M.C. van Bohemen, J. Sietsma, Modeling of isothermal bainite formation based on the nucleation kinetics, Int. J. Mater. Res. 99 (7) (2008) 739–747, http://dx.doi.org/10.3139/146.101695.

[39] F. Mühl, Simulationsgestützte Prozessoptimierung des Wärmebehandlungsverfahrens Internal Quenching (Ph.D. thesis), Karlsruher Institut für Technologie (KIT), Karlsruhe, 2022, http://dx.doi.org/10.5445/IR/1000145728.

[40] S. van Bohemen, J. Sietsma, Martensite formation in partially and fully austenitic plain carbon steels, Met. Mater. Trans. A 40 (5) (2009) 1059–1068, http://dx.doi.org/10.1007/s11661-009-9796-2.

[41] D. Kaiser, B. Graaff, S. Dietrich, V. Schulze, Investigation of the precipitation kinetics and microstructure evolution of martensitic AISI 4140 steel during tempering with high heating rates, Met. Res. Technol. 115 (4) (2018) 404, http://dx.doi.org/10.1051/metal/2018026.

[42] P. Peyre, P. Aubry, R. Fabbro, R. Neveu, A. Longuet, Analytical and numerical modelling of the direct metal deposition laser process, J. Phys. D: Appl. Phys. 41 (2) (2008) 025403, http://dx.doi.org/10.1088/0022-3727/41/2/025403.

[43] P. Schüßler, V. Schulze, S. Dietrich, Code repository "3mf-to-pbflb-gcode", 2026, http://dx.doi.org/10.35097/C5R9SSKM0PAN3GGK.

[44] F.A. Gers, J. Schmidhuber, F. Cummins, Learning to forget: Continual prediction with LSTM, Neural Comput. 12 (10) (2000) 2451–2471, http://dx.doi.org/10.1162/089976600300015015.

[45] R.J. Williams, D. Zipser, A learning algorithm for continually running fully recurrent neural networks, Neural Comput. 1 (2) (1989) 270–280, http://dx.doi.org/10.1162/neco.1989.1.2.270.

[46] P. Schüßler, V. Schulze, S. Dietrich, Training-validation-testing dataset for "real-time prediction of thermal history and hardness in laser powder bed fusion using deep learning", 2026, http://dx.doi.org/10.35097/pmem1cb9gu1ck8xz.

[47] P. Schüßler, V. Schulze, S. Dietrich, Trained LSTM ensemble models for "real-time prediction of thermal history and hardness in laser powder bed fusion using deep learning", 2026, http://dx.doi.org/10.35097/37da9d66y4t27q55.

[48] P. Schüßler, V. Schulze, S. Dietrich, Code repository for "real-time prediction of thermal history and hardness in laser powder bed fusion using deep learning", 2026, http://dx.doi.org/10.35097/dg39f4p0wxqdnfxy.