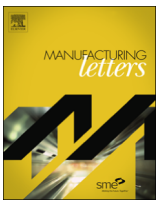
Available online at www.sciencedirect.com

Manufacturing Letters

Manufacturing Letters 49 (2026) 204–215



54th SME North American Manufacturing Research Conference (NAMRC 54, 2026)

UNiCO-AM: Uncertainty-aware Neural Co-design for Additive Manufacturing via an Attention-based Spatio-Temporal Neural Operator and Bayesian Optimization

Vispi Nevile Karkaria[a], Ying-Kuan Tsai[a], Yi-Ping Chen[a], Dex Doksoo Lee[a], Jian Cao[a], Wei Chen[a]

[a]*Department of Mechanical Engineering, Northwestern University, Evanston, Illinois, USA*

Abstract

Metal Additive Manufacturing (AM) is a complex process where the interaction between material and process parameters can significantly influence melt-pool dynamics, thermal gradients, and solidification behavior, ultimately determining the resulting microstructure, mechanical properties, and overall quality of the fabricated part. However, inherent challenges such as the high-dimensional process-material design space, strong nonlinearity in thermal-physical interactions, and the high computational cost of high-fidelity simulations make it difficult to systematically explore and optimize these interdependencies across diverse alloys and process conditions. To address this challenge, we propose *UNiCO-AM* (Uncertainty-aware Neural Co-design for Additive Manufacturing), a physics-informed optimization framework that integrates an attention-based spatio-temporal neural operator (ASNO) with Bayesian optimization (BO). In this framework, we develop a material–process-agnostic ASNO model that is capable of learning the coupled spatio-temporal thermodynamics of the additive manufacturing process, accurately predicting full-field temperature evolution across different alloys and laser parameters, and generalizing to unseen material–process combinations. Another key contribution of this work is the quantification of epistemic uncertainty in ASNO using a Linear Laplace-based Approximation (LLA) method. This approach estimates the model's uncertainty around its learned parameters, providing the confidence information required for effective BO. Incorporating uncertainty quantification not only enables more reliable and data-efficient optimization but also enhances the generalization capability of ASNO across diverse materials and process conditions. Finally, we complete the *UNiCO-AM* framework by establishing a concurrent material-process BO loop for AM, incorporating a q-Expected Hypervolume Improvement (q-EHVI) strategy to efficiently explore the high-dimensional design space under physical and thermal constraints. The primary objective of the *UNiCO-AM* framework is to enhance the strength of the fabricated part by promoting finer microstructural grain sizes through joint optimization of process parameters and material properties that regulate cooling rates and solidification behavior during AM. This approach yields a Pareto front of material–process combinations that balance cooling efficiency and thermal stability, leading to improved mechanical performance in fabricated parts while also enabling the discovery of novel material-process property combinations.

Keywords: Additive manufacturing; concurrent design; Bayesian optimization; neural operators; uncertainty quantification; thermal modeling

1. Introduction

Metal additive manufacturing (AM) is becoming an important method for producing lightweight, complex, and high-performance components used in aerospace, energy, and biomedical systems. Reliable part quality in AM requires precise control of the thermal history during fabrication. Such control depends on both process and material optimization, which together enable consistent and defect-free builds [1, 2]. The thermal response of a build results from the combined effects of alloy-specific material properties and laser-driven heating, which govern melt pool behavior and determine the final microstructure and performance of the part [3, 4]. In directed energy deposition (DED), changes in laser power, scan speed, or hatch spacing strongly affect the cooling rate and temperature gradients, leading to different solidification patterns and microstructural evolution [5, 6].

As AM evolves from empirical parameter tuning to systematic, model-driven design, optimization must move beyond isolated process settings to jointly account for material behavior. Such integration enables the discovery of process–material combinations that tailor the microstructure of the fabricated

* Wei Chen
E-mail address: weichen@northwestern.edu (Wei Chen).

part, which is primarily governed by the cooling rate and melt-pool morphology [7]. Recent research has increasingly recognized that process parameters alone cannot guarantee manufacturability or performance, motivating *co-design* formulations that integrate materials or structural designs into process optimization rather than treating them as sequential or decoupled tasks [8]. For example, structure-process co-design has been explored to simultaneously tailor part topology and scan strategies, yielding designs that remain printable while improving mechanical performance [9]. In the topology optimization considering AM processes, researchers have integrated thermo-mechanical constraints into topology optimization problems to mitigate warping, residual stress, and distortion during metal AM builds [10]. Another example is imposing overheating or melt-pool stability constraints to ensure feasible thermal fields while maximizing structural efficiency [11]. Together, these studies demonstrate the value of embedding physics-based process models into design loops. However, most existing frameworks still assume fixed material properties or restrict analysis to a single alloy system, which limits their ability to generalize and prevents the discovery of new material–process combinations that could expand the feasible design space. Moreover, these approaches rarely incorporate high-fidelity, spatio-temporal evaluations of melt pool dynamics, leading to inaccurate thermal predictions and suboptimal process control when applied to new materials.

Building upon the growing need for concurrent process-material optimization in AM, several fundamental challenges still hinder the realization of scalable concurrent design in AM [12]. First, the decision space is inherently high-dimensional and fully continuous, encompassing not only multi-parameter laser power profiles (e.g., power amplitude, pulse duration, frequency modulation, and dwell scheduling) but also continuously varying temperature-dependent material properties such as thermal conductivity, heat capacity, density, and absorptivity [13]. These factors interact nonlinearly with transient melt-pool evolution in ways that are highly sensitive to local thermal history, making the design landscape strongly nonconvex for an optimization method [14]. As each high-fidelity DED simulation is computationally expensive, multi-objective optimization must proceed with a minimal number of evaluations, making Bayesian optimization (BO) a suitable candidate [15]. However, conventional Gaussian Process (GP)-based BO methods scale poorly with dimensionality and struggle to capture the complex physics-governed interactions across time and space, necessitating surrogate models such as neural operators that can efficiently embed high-dimensional relationships in functional form [16]. Moreover, existing scientific machine learning (SciML) surrogates often struggle to represent the coupled spatio-temporal dynamics of AM thermal fields across diverse material–process combinations, particularly under rapid heating-cooling cycles, steep temperature gradients, and dynamic melt-pool boundary motion [17].

To overcome these limitations, recent advances have focused on developing SciML paradigms, particularly neural operators and emerging foundation models, that enable generalizable estimation of spatio-temporal behaviors for unseen materials by learning shared latent representations across alloys and process regimes. Such architectures disentangle temporal progression from spatial redistribution [18]. Finally, practical deployment in industrial AM benefits from batch optimization strategies that can evaluate multiple promising candidates per iteration, both to reduce simulation costs and to identify Pareto trade-offs [19]. Without batch the optimization might become prohibitively slow and may converge to narrow local optima that fail to capture the broader material-dependent behavior.

To address these gaps, we introduce UNiCO-AM, an uncertainty-aware neural co-design framework that unifies spatio-temporal surrogate learning with multi-objective, batch BO for concurrent process–material design. UNiCO-AM's machine learning model learns the spatio-temporal behavior that is inherent to additive manufacturing due to the continuous movement of the heat source, localized melting and solidification, and dynamic thermal interactions between successive layers and tracks [**?**]. For quantifying epistemic uncertainty within the surrogate, we employ the Linear Laplace Approximation (LLA) [20]. Finally UNiCO-AM utilizes the multi-objective BO which aids in achieving concurrent material-process optimization. The following contributions are made in this paper:

- A unified co-design framework named UNiCO-AM is developed for optimizing high-dimensional, fully continuous process and material properties in additive manufacturing by utilizing a neural operator-based surrogate within a multi-objective BO framework.
- A material-process agnostic attention-based spatio-temporal neural operator is developed for modeling of AM thermal evolution, with Laplace-based uncertainty quantification enabling uncertainty-aware prediction.
- A batch-based, multi-objective BO pipeline in *UNiCO-AM* is developed that efficiently generates Pareto-optimal process-material combinations that improves the material strength of the fabricated part.

These contributions motivate a clear and practical framework that treats materials consistently across varying thermophysical domains, scales to high-dimensional continuous design spaces, and supports uncertainty-aware decision-making. While UNiCO-AM is formulated for offline concurrent process–material optimization, the trained ASNO surrogate enables millisecond-level melt-pool thermal predictions that can naturally support near real-time supervisory control and future integration within receding-horizon or model predictive control (MPC) feedback loops. Section 2 first introduces the DED AM simulation model, *GAMMA*, which forms the foundation for generating high-fidelity thermal fields and melt-pool dynamics. It then details the theoretical formulation of thermal physics, the material-normalized cooling objectives, and the associated feasibility constraints, leading to the formal formulation of the optimization problem and its uncertainty quantification methodology. Section 3 describes the UNiCO-AM methodology, including the scientific machine learning surrogate ASNO, the Laplace based uncertainty quantification procedure, and the Bayesian optimization strategy for concurrent process and material design. Section 4 presents results, first validating surrogate accuracy and cross material generalization and then evaluating Bayesian optimization performance. Section 5 concludes

with implications for concurrent process and material design and outlines practical directions for deployment.

2. Theoretical Framework

Building on the problem formulation, we establish the theoretical foundation for process-material co-design in AM. We first describe the DED simulation model governing transient heat transfer and melt-pool evolution, then define cooling rate as a thermal performance metric. A material-normalization framework is introduced to compare thermal responses across alloys, followed by an uncertainty-aware optimization formulation that integrates these metrics into the surrogate and multi-objective BO framework.

2.1. Thermal Physics of AM and Cooling Rate

In the UNiCO-AM framework, a high-fidelity simulation model is essential for training and updating the surrogate within the BO. This paper employs the *GAMMA* simulation [12], which models conduction-dominated transient heat-transfer behavior driving the thermal response in directed energy deposition (DED), where laser–material interaction and temperature-dependent properties determine melt-pool formation, evolution, and solidification. *GAMMA* resolves thermally driven conduction physics and does not explicitly model multiphase fluid effects such as keyholing, vapor depression, or other melt-pool defect mechanisms, and the present work therefore assumes operation within a stable, non-defect regime. These thermal dynamics strongly influence microstructure formation, residual stress development, and overall part quality [21, 22, 23]. In laser-based DED, heat transfer is mainly governed by conduction with localized, time-varying laser energy input coupled to temperature-dependent material properties, and the transient heat equation representing this behavior over a spatial domain $\Omega \subset \mathbb{R}^3$ is expressed as

$$\rho(T)\, C_p(T)\, \frac{\partial T(\mathbf{X},t)}{\partial t} = \nabla \cdot [k(T)\, \nabla T(\mathbf{X},t)] \; + \; Q(\mathbf{X},t), \quad (1)$$

where ρ is the material density, C_p is the specific heat capacity, k is the thermal conductivity, T is temperature, Q denotes the volumetric laser energy input, and $\mathbf{X} = (X, Y, Z)$ represents the spatial coordinate vector within the computational domain.

Boundary effects such as convection and radiation at free surfaces, as well as powder-solid interfacial continuity, may be considered but are typically secondary during the short dwell times characteristic of DED [24]. The laser heat source Q is commonly modeled as a moving Gaussian distribution characterized by absorptivity η, beam radius r_b, and scan velocity v, centered along the laser path $\mathbf{X}_\ell(t)$. This formulation reproduces classical moving-source behavior in the constant-property limit. As the moving heat source governs the transient temperature gradients and solidification dynamics, it directly determines the thermal metrics that characterize process performance, most notably the cooling rate.

Among the thermally driven process indicators in AM that improve material properties, the cooling rate plays a pivotal role in determining the quality and performance of fabricated components [25, 26, 27]. It governs the rate at which the melt pool solidifies, thereby influencing grain morphology, microstructure refinement, and defect formation, such as porosity or cracking. Faster cooling rates generally promote fine-grained, homogeneous microstructures with better mechanical properties, whereas excessively high rates may introduce thermal stresses or incomplete fusion. Hence, planning of the cooling rate is essential for achieving a balance between microstructural refinement and process stability. To quantitatively assess this phenomenon, we define the instantaneous or local cooling rate as

$$\dot{T}(\mathbf{X},t) \equiv \frac{\partial T(\mathbf{X},t)}{\partial t}, \quad (2)$$

which characterizes the temporal temperature decay at each spatial location $\mathbf{X}$ and time t within the melt pool. The spatial and temporal distribution of $\dot{T}$ provides rich information about the thermal history of the part and serves as a direct proxy for solidification kinetics. In this work, cooling rate is evaluated specifically within the alloy-dependent solidification interval, i.e., for $T_S \le T(\mathbf{X},t) \le T_L$, where T_S and T_L denote the solidus and liquidus temperatures, respectively. To isolate this temperature range, we introduce a masking function $\chi(\mathbf{X},t)$ that identifies the actively solidifying region:

$$\chi(\mathbf{X},t) = \begin{cases} 1, & \text{if } T_S \le T(\mathbf{X},t) \le T_L, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

A scalar metric suitable for optimization is then derived by spatially averaging $\dot{T}$ over the melt-pool region where $\chi(\mathbf{X},t) = 1$, followed by temporal aggregation over the solidification duration Δt. Denoting spatial and temporal averaging as $\langle\cdot\rangle_{\Omega_{\text{ROI}}}$ and $\langle\cdot\rangle_{\Delta t}$, respectively, the effective cooling rate is computed as

$$\dot{T}_{\text{eff}} \;=\; \Big\langle \big\langle \dot{T}(\mathbf{X},t)\,\chi(\mathbf{X},t) \big\rangle_{\Omega_{\text{ROI}}} \Big\rangle_{\Delta t}. \quad (4)$$

This effective cooling rate serves as a compact yet physically meaningful indicator of thermal performance in DED.

2.2. Material-Normalized Cooling Objective

While the previous section established cooling rate as a key thermal indicator of solidification and defect formation, its interpretation remains inherently material-dependent. Alloys differ in thermophysical properties and solidification intervals, so identical process conditions can yield markedly different cooling responses. For instance, alloys with higher thermal diffusivity or narrower freezing ranges cool faster even under the same laser parameters. Directly maximizing the raw cooling rate $\dot{T}_{\text{eff}}$ would therefore bias optimization toward highly conductive alloys rather than identifying those performing efficiently relative to their intrinsic physical limits.

To enable consistent, physics-informed comparison across alloys, we introduce a material-normalized cooling metric that incorporates characteristic conductive and thermal scaling. This normalization transforms the absolute cooling rate into a dimensionless measure of thermal efficiency, quantifying how effectively a process-material pair dissipates heat relative to its intrinsic thermophysical properties. Let L denote the characteristic thermal length scale (the melt-pool half-width), $\alpha(T) = k(T)/[\rho(T)\, C_p(T)]$ represent the temperature-dependent thermal diffusivity, and $\Delta T_m = T_L - T_S$ the alloy-specific solidification interval. The normalized cooling rate is then defined as

$$\widehat{\dot{T}}_{\text{eff}}(\mathbf{X}, t) = \frac{\dot{T}(\mathbf{X}, t)\, L^2}{\alpha^\star\, \Delta T_m}, \tag{5}$$

scaling $\dot{T}$ by conductive and thermal intervals to quantify process-material cooling efficiency [28, 25]. After spatial and temporal aggregation, we obtain the effective normalized cooling rate $\widehat{\dot{T}}_{\text{eff}}$.

For statistical comparability across alloys, the effective rate is standardized using a z-score:

$$\dot{T}_{\text{rel}} = \frac{\widehat{\dot{T}}_{\text{eff}} - \mu_m}{\sigma_m}, \tag{6}$$

where μ_m and σ_m denote the mean and standard deviation of $\widehat{\dot{T}}_{\text{eff}}$ for material m obtained from prior data. This normalization aligns performance regions across materials and mitigates sampling bias in optimization.

Since consistent cooling is as critical as rapid cooling as a high cooling rate in one spatial region compared to another can lead to microstructural inconsistencies resulting in a weaker build. Therefore, we further quantify variability through the coefficient of variation:

$$CV_{\widehat{\dot{T}}} = \frac{\sigma_{\widehat{\dot{T}}}}{\mu_{\widehat{\dot{T}}}}, \tag{7}$$

where $\sigma_{\widehat{\dot{T}}}$ and $\mu_{\widehat{\dot{T}}}$ represent the predictive standard deviation and mean of normalized cooling. Accordingly, the co-design problem seeks to (i) maximize $\widehat{\dot{T}}_{\text{eff}}$ for rapid, uniform solidification and (ii) minimize $CV_{\widehat{\dot{T}}}$ for thermal stability. These competing goals define a multi-objective optimization:

$$\max_{\mathbf{x}} \left[\widehat{\dot{T}}_{\text{eff}}(\mathbf{x}),\ -CV_{\widehat{\dot{T}}}(\mathbf{x})\right], \tag{8}$$

where $\mathbf{x}$ represents the design variables, which will be detailed in Section 2.4. The solution set is called the Pareto front, which consists of non-dominated solutions representing optimal trade-offs among conflicting objectives. The Pareto front identifies thermally optimal process-material combinations balancing cooling intensity and stability.

These two objectives represent a robust design, aiming to enhance performance while minimizing sensitivity to process-material variations. Maximizing the cooling rate promotes rapid solidification, whereas minimizing its coefficient of variation ensures uniform thermal behavior, reducing defects from stochastic melt-pool dynamics.

2.3. Constraints

Achieving concurrent objectives requires process-material combinations to remain physically feasible and manufacturable. In DED, exceeding thermal or geometric limits can cause over-melting or instability, so analytical constraints are embedded to maintain physical limits and ensure generality across materials. In this paper, we define three constraints to have stable and manufacturable builds.

The first constraint is defined as a penetration-depth limit on the melt-pool geometry to prevent substrate erosion and ensure stable layer formation, and it is defined as:

$$g_1(\mathbf{x}) = D(\mathbf{x}) - D_{\max} \le 0, \tag{9}$$

where $D_{\max}$ denotes the maximum allowable penetration depth calculated from high-fidelity DED studies. A representative upper bound of $D_{\max} = 0.15\,\text{mm}$ is adopted for all alloys, corresponding to the threshold beyond which excessive substrate erosion typically occurs [29].

In the second constraint, peak temperature is restricted to reduce residual stress accumulation:

$$g_2(\mathbf{x}) = T_{\max}(\mathbf{x}) - T_{\text{fail}} \le 0, \tag{10}$$

where T_{fail} represents the onset of vaporization or structural instability. A general limit of $T_{\text{fail}} = 2173\,\text{K}$ is selected as a conservative upper bound for metallic alloys in DED [30].

Finally, in the third and fourth constraints, temperature-dependent thermophysical properties are required to vary smoothly within the solidification interval to maintain physical consistency and numerical stability.

$$g_3(\mathbf{x}) = \max_{T\in[T_S,T_L]} \left[-\frac{dk(T)}{dT}\right] \le 0,\ g_4(\mathbf{x}) = \max_{T\in[T_S,T_L]} \left[-\frac{dC_p(T)}{dT}\right] \le 0, \tag{11}$$

which enforces monotonic (non-decreasing) behavior of conductivity $k(T)$ and heat capacity $C_p(T)$ over the range $T_S \le T \le T_L$. Adding these constraints ensures that optimization follows thermophysical laws and manufacturability.

2.4. Optimization Problem and Uncertainty Handling

Building upon the objectives and physical constraints introduced in the previous section, we now formalize the concurrent process–material co-design as a Bayesian optimization problem. The design variables are denoted as $\mathbf{x} = [\mathbf{x}_p^\top, \mathbf{x}_m^\top]^\top$, representing a joint vector of process and material parameters. The process variables $\mathbf{x}_p$ correspond to a parametric laser power profile [$\mathbf{P}_{\text{laser}}$], generated using eleven parameters that control amplitude, frequency, phase, rates of change, and temporal modulation [4], thereby capturing the full range of dynamic energy-input strategies used in AM. The material variables $\mathbf{x}_m$ consist of [ρ_{density}, T_{solidus}, T_{liquidus}, $\boldsymbol{C}_p(T)$, $\boldsymbol{k}(T)$], where the scalar quantities describe bulk density and phase-transition temperatures, while the temperature-dependent vectors $\boldsymbol{C}_p(T)$ and $\boldsymbol{k}(T)$ are discretized at seven representative temperature points to capture the shape and curvature of heat-capacity and thermal-conductivity profiles. Seven representative temperature points are selected as an effective and physically meaningful resolution because melt-pool evolution is governed by strongly temperature-dependent diffusion physics in Eq. (1), where both $k(T)$ and $C_p(T)$ vary nonlinearly across the solidification range $[T_S, T_L]$ and directly modulate thermal diffusivity $\alpha(T) = k(T)/(\rho C_p)$. In particular, accurately resolving property variation in this interval is critical for capturing latent-heat-dominated enthalpy buffering near phase transition, conductivity-driven heat extraction at high temperatures, and the sharp gradients that control cooling-rate stability and melt-pool morphology. This discretization therefore provides sufficient functional expressivity to represent physically realizable thermophysical curvature while avoiding unnecessary dimensionality that would degrade BO efficiency. A sensitivity study with coarser (5-point) and finer (9-point) discretizations further confirmed that seven temperature points yield stable surrogate

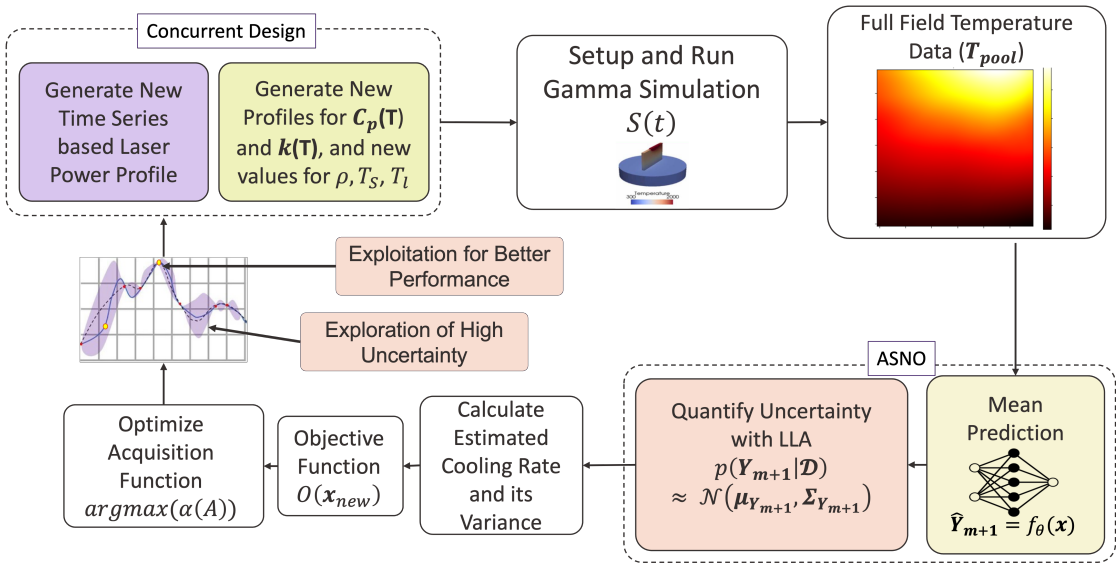


Figure 1. *UNiCO-AM* (Uncertainty-aware Neural Co-design for Additive Manufacturing) Framework

accuracy and consistent Pareto-optimal optimization trends, supporting the robustness of the chosen parameterization. This parameterization enables a continuous material design space in which the optimizer can search for new combinations of effective thermophysical properties that may not correspond to a single known alloy but nevertheless reflect physically realizable behaviors. Such a formulation naturally supports inverse design: instead of selecting from discrete alloys, the BO loop directly optimizes the material-property functions that govern melt-pool dynamics, allowing the discovery of novel or hybrid material behaviors that produce desirable cooling rates and microstructures. In total, the 28-dimensional process–material design vector captures the most influential thermophysical and energy-input factors controlling melt-pool evolution, solidification pathways, and resulting part integrity in metal additive manufacturing.

With the complete set of process-material design variables defined, the co-design task aims to identify optimal combinations of design variables that jointly maximize cooling rate and minimize variability. The two competing objectives are expressed as:

$$\max_{\mathbf{x}} \ [\,\mu_{\widehat{T}}(\mathbf{x}),\ -\mu_{CV}(\mathbf{x})\,], \tag{12}$$

where $\mu_{\widehat{T}}(\mathbf{x})$ denotes the expected normalized cooling rate and $\mu_{CV}(\mathbf{x})$ the expected coefficient of variation.

To ensure feasibility, each physics-based constraint from the previous section is evaluated probabilistically through chance constraints [31, 32]:

$$\Pr(g_i(\mathbf{x}) \le 0) \ge 1 - \epsilon_i, \quad i = 1, 2, 3, 4, \tag{13}$$

which, under the Gaussian approximation $g_i(\mathbf{x}) \sim \mathcal{N}(\mu_{g_i}(\mathbf{x}), \sigma^2_{g_i}(\mathbf{x}))$, reduces to

$$\mu_{g_i}(\mathbf{x}) + \Phi^{-1}(1 - \epsilon_i)\sigma_{g_i}(\mathbf{x}) \le 0, \quad i = 1, 2, 3, 4, \tag{14}$$

where Φ^{-1} is the standard normal quantile function and $\epsilon_i = 10^{-3}$ represents the tolerance for constraint violations, where smaller values indicate a more conservative treatment of constraints. In this work, the tolerance is chosen to balance reliability and design flexibility in optimizing the material properties and process conditions for the DED process.

The complete uncertainty-aware multi-objective optimization problem is therefore formulated as

$$\begin{aligned} \max_{\mathbf{x}} \quad & \left[\mu_{\widehat{T}}(\mathbf{x}),\ -\mu_{CV}(\mathbf{x})\right] \\ \text{s.t.} \quad & \mu_{g_i}(\mathbf{x}) + \Phi^{-1}(1 - \epsilon_i)\sigma_{g_i}(\mathbf{x}) \le 0, \quad i = 1, 2, 3, 4. \end{aligned} \tag{15}$$

This formulation ensures that Pareto-optimal solutions reflect a balance between cooling efficiency, thermal stability, and physical feasibility. It lays the foundation for the BO framework of UNiCO-AM.

3. Methodology

In this section, we introduce *UNiCO-AM*, an uncertainty-aware neural co-design framework tailored for metal additive manufacturing (AM). The framework enables high-dimensional concurrent material-process optimization by bridging physics-based simulation and temporal-spatial machine learning to enable scalable, data-efficient exploration of the design space. As shown in Figure 1, *UNiCO-AM* integrates three synergistic components within a closed-loop optimization workflow. The first component is the Attention-Based Spatio-Temporal Neural Operator (ASNO) [33], which is as a high-fidelity surrogate for predicting melt-pool temperature evolution under varying alloy compositions and laser power process conditions. ASNO captures the coupled spatio-temporal thermophysical behavior governing melt-pool formation, enabling accurate prediction of transient thermal fields, cooling rates, and solidification dynamics that control microstructural evolution in AM builds. To quantify the uncertainty of these predictions, a Linear Laplace-based Approximation (LLA) is employed to estimate epistemic uncertainty in ASNO's outputs, providing confidence intervals of meltpool.

The second component of *UNiCO-AM* couples these uncertainty-calibrated predictions with a multi-objective Bayesian optimization (BO) framework designed for concurrent material–process optimization in AM. The BO module

leverages ASNO's mean and variance predictions to efficiently navigate the high-dimensional design space, where experimental or simulation-based evaluations are costly. Using a q-Expected Hypervolume Improvement (q-EHVI) acquisition strategy, *UNiCO-AM* identifies Pareto-optimal trade-offs between key thermal metrics by maximizing effective cooling rate while minimizing its variability, subject to manufacturability and physical feasibility constraints. Through this iterative process of prediction, uncertainty quantification, and optimization, *UNiCO-AM* establishes a robust and physics-informed co-design paradigm that accelerates the discovery of optimal material–process combinations, enhances microstructural control, and reduces dependence on computationally intensive thermal simulations in metal additive manufacturing.

3.1. Scientific Machine Learning Surrogate to Predict Melt Pool Window

A core component of UNiCO-AM is a physics-aware machine learning surrogate that bridges high-fidelity DED simulations with BO. Although ASNO does not impose PDE residual constraints as in classical PINNs, it is physics-informed in the operator-learning sense because it learns the solution evolution of the governing transient heat-transfer equation (Eq. (1)) from high-fidelity GAMMA simulations while explicitly conditioning on thermophysical inputs such as $k(T)$ and $C_p(T)$ that control thermal diffusion. The surrogate employs the Attention-based Spatio-Temporal Neural Operator (ASNO), a Neural Operator architecture that learns mappings between input and output function spaces, enabling efficient modeling of field-to-field relationships governed by partial differential equations. This formulation allows ASNO to accurately capture transient heat transfer across alloys and process conditions. In this work, ASNO predicts the *melt-pool window*, a localized spatio-temporal region centered on the melt pool that records the temperature distribution and its temporal evolution. As shown in Figure 2, a moving melt-pool window of size 10 mm × 10 mm continuously tracks the laser position, capturing the dynamic evolution of the thermal field. Each window is discretized into a 30 × 30 spatial grid, resulting in 900 points that represent the temperature field in the active region. By focusing on this localized, physically dominant zone, ASNO effectively models spatio-temporal temperature evolution, accurately reproducing heat transfer, cooling rate, and solidification behavior across diverse process-material combinations.

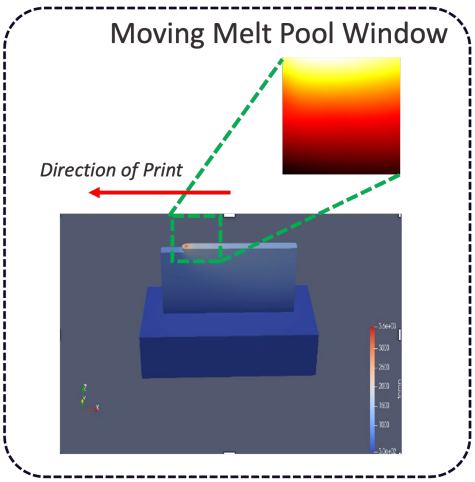


Figure 2. Moving melt-pool window in ASNO capturing transient temperature evolution along the print path.

Building on this melt pool window, ASNO, parameterized by $\boldsymbol{\theta} = [\,\boldsymbol{\theta}_{\text{TE}}, \boldsymbol{\theta}_{\text{NAO}}\,]$, decouples time evolution from spatial correction to improve generalization. ASNO comprises two steps, namely the Time-Series Temporal Encoder (TE) and the Spatial Non-Local Attention Operator (NAO), where the TE captures short- and long-term temporal dependencies in the evolving thermal sequence, and the NAO corrects spatial consistency by modeling process-material interactions within each melt-pool window. The first step, the TE, as shown in Figure 3, takes in a sequence of previous melt-pool windows $\mathbf{Y}_{m-n+1:m}$ to predict the next latent melt-pool temperature window, effectively capturing the evolving local dynamics of the deposition process of DED:

$$\mathbf{H}_{m+1} = \text{TE}(\mathbf{Y}_{m-n+1:m};\, \boldsymbol{\theta}_{\text{TE}}), \tag{16}$$

The output of the TE forms a latent melt-pool window representation that forecasts the temporal evolution of local thermal features. This latent representation is then passed to the NAO [33] along with the joint process-material inputs $\mathbf{x} = [\mathbf{x}_m, \mathbf{x}_p]$, enabling the model to correct spatial correlations and incorporate process-driven and material-dependent effects before generating the future predicted melt-pool window:

$$\begin{aligned}\widehat{\mathbf{Y}}_{m+1} &= \text{NAO}(\mathbf{H}_{m+1};\, \mathbf{X},\, \boldsymbol{\theta}_{\text{NAO}}) \\ &= \mathbf{H}_{m+1} + \int_{\Omega} \boldsymbol{\alpha}(x, x';\, \mathbf{X},\, \boldsymbol{\theta}_{\text{NAO}})\, W_{\boldsymbol{\phi}}(\mathbf{H}_{m+1}(x'))\, dx',\end{aligned} \tag{17}$$

where α is a learned attention kernel encoding long-range, process-material coupling. Together, these operators define the forward ASNO prediction of melt pool window:

$$\begin{aligned}\widehat{\boldsymbol{Y}}_{m+1} &= \text{ASNO}\,(\boldsymbol{Y}_{m-n+1:m},\, \boldsymbol{X};\, \boldsymbol{\theta}) \\ &= \text{NAO}\,(\text{TE}\,(\boldsymbol{Y}_{m-n+1:m};\, \boldsymbol{\theta}_{\text{TE}})\,,\, \boldsymbol{X},\, \boldsymbol{\theta}_{\text{NAO}})\,.\end{aligned} \tag{18}$$

This coupling ensures that ASNO captures both the short-term temporal dependencies and the long-range spatial correlations inherent to melt-pool evolution, resulting in physically consistent full-field prediction. The model parameters $\boldsymbol{\theta}$ are learned via MAP estimation using a mean squared error (MSE) loss, where the resulting MAP solution $\boldsymbol{\theta}^{\star}$ initializes uncertainty quantification.

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{j=1}^{N}\frac{1}{T_j}\sum_{m}\left\|\widehat{\boldsymbol{Y}}^{(j)}_{m+1} - \boldsymbol{Y}^{(j)}_{m+1}\right\|_2^2. \tag{19}$$

For training, a short-to-long rollout curriculum (10, 25, and 50 time steps) with a stratified 80–10–10 split across process-material pairs is used to ensure balanced temporal and material coverage. Low-rank adapters with rank $r = 8$ are employed within the Transformer layers to enable rapid fine-tuning and efficient adaptation to new alloys using as few as 50 additional samples. The ASNO surrogate is trained on high-fidelity melt-pool temperature data generated using the *GAMMA* DED simulator, which provides temperature-dependent thermal fields across five representative alloys: 316L stainless steel, 1018 Steel, IN625, Ti-6Al-4V, and IN718. The material properties is given in this document [34]. In this paper, all the analysis

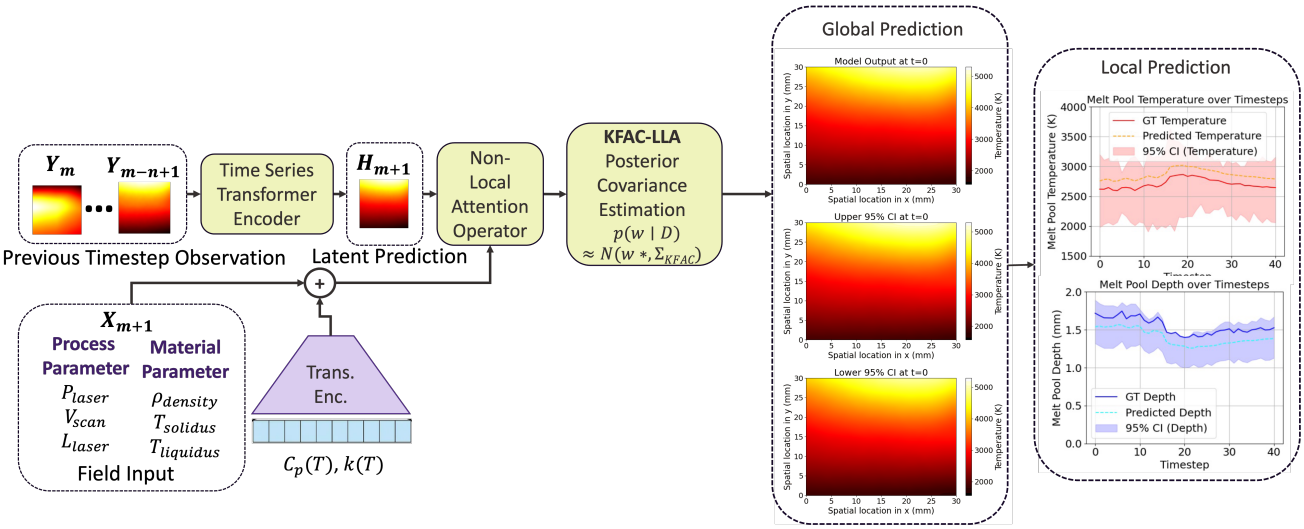


Figure 3. ASNO architecture showing melt-pool input, spatio-temporal prediction, and uncertainty estimation.

is done on a thin-walled part. For training 100 unique laser power profiles are generated for each material. For each profile, melt-pool temperature windows are recorded for 720 s with a time step of 0.2 s, corresponding to 3600 temporal samples per trajectory. The complete dataset therefore comprises 1,800,000 melt-pool temperature windows, of which 1,440,000 splits to the training dataset (400 profiles) and 360,000 to the testing dataset (100 profiles), following an 80%–20% division stratified across both process and material space to ensure balanced training. During training, data are organized into batches of 64 samples. Each training input includes five preceding melt-pool temperature windows, denoted as $\boldsymbol{Y}_{m-n+1:m}$, which capture the recent spatio-temporal evolution of the melt-pool region. Thermal inputs include process-material inputs $\boldsymbol{x} = [\boldsymbol{x}_p, \boldsymbol{x}_m]$, where $\boldsymbol{x}_p = [\boldsymbol{P}_{\text{laser}}]$, and $\boldsymbol{x}_m = [\rho_{\text{density}}, T_{\text{solidus}}, T_{\text{liquidus}}, \boldsymbol{C}_p(T), \boldsymbol{k}(T)]$. The model is trained for 500 epochs with a learning rate scheduler using a decay factor of $\gamma = 0.1$, ensuring stable convergence and robust generalization across alloys. All experiments were conducted on a single NVIDIA RTX A6000 GPU, and once trained ASNO achieves millisecond-level inference per melt-pool window, supporting scalable evaluation of larger AM builds through localized operator prediction. Once trained, the network predicts the next melt-pool temperature window $\widehat{\boldsymbol{Y}}_{m+1}$, enabling autoregressive rollout of melt-pool temperature window prediction during optimization.

3.2. Uncertainty Quantification via Laplace Approximation

In this section, we introduce the epistemic uncertainty quantification used in ASNO. In the BO context, quantification of epistemic uncertainty is vital for guiding exploration and preventing overconfident extrapolation in data-scarce regions. Epistemic uncertainty stems from the model's limited knowledge of the true data distribution and can be reduced as more melt-pool temperature observations are collected. To address this, ASNO employs the LLA, which treats the trained network as locally linear around its MAP parameters and approximates the posterior as a Gaussian. This method calculates the standard deviation of model parameters and propagates parameter uncertainty through the network to capture how small weight perturbations affect predictions, thereby enabling uncertainty-aware outputs without expensive sampling. The parameter standard deviation is then obtained via MAP estimation:

$$\boldsymbol{\theta}_{\text{MAP}} = \arg\min_{\boldsymbol{\theta}} \left[\ell(\mathcal{D}, \boldsymbol{\theta}) + r(\boldsymbol{\theta})\right], \tag{20a}$$

$$\boldsymbol{\theta} \sim \mathcal{N}(\boldsymbol{\theta}_{\text{MAP}}, \boldsymbol{\Sigma}), \qquad \boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}, \tag{20b}$$

$$\boldsymbol{\Lambda} = \nabla^2_{\boldsymbol{\theta}}[\ell(\mathcal{D}, \boldsymbol{\theta}) + r(\boldsymbol{\theta})]\big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{\text{MAP}}}, \tag{20c}$$

$$\boldsymbol{\Sigma} \approx \boldsymbol{\Lambda}^{-1} \approx \text{diag}\Big(\tfrac{1}{\lambda_1}, \tfrac{1}{\lambda_2}, \ldots, \tfrac{1}{\lambda_n}\Big), \tag{20d}$$

These equations collectively define the LLA framework used in ASNO to quantify epistemic uncertainty. Equation (20a) determines the MAP estimate $\boldsymbol{\theta}_{\text{MAP}}$, representing the most likely network parameters given the training data $\mathcal{D}$ and prior knowledge. Equation (20b) models the posterior distribution of the parameters as a Gaussian centered at $\boldsymbol{\theta}_{\text{MAP}}$, with covariance $\boldsymbol{\Sigma}$ equal to the inverse of the curvature matrix $\boldsymbol{\Lambda}$. Equation (20c) defines $\boldsymbol{\Lambda}$ as the second-order derivative (curvature) of the regularized loss landscape evaluated on $\mathcal{D}$, and Equation (20d) introduces a diagonal approximation to make computation tractable for high-dimensional neural operators. This simplification assumes parameter independence, greatly reducing computational cost while preserving meaningful curvature information. As a result, ASNO can efficiently propagate epistemic uncertainty through its predictive model of melt-pool temperature windows.

With the model parameter covariance $\boldsymbol{\Sigma}$, uncertainty is propagated through ASNO's predictive map. Linearizing Eq. (18) around $\boldsymbol{\theta}_{\text{MAP}}$ gives the predictive covariance and its associated Jacobian as follows:

$$\text{Cov}[\boldsymbol{Y}_{m+1}] \approx \boldsymbol{J}_{\boldsymbol{\theta}}^{(m+1)} \boldsymbol{\Sigma} \left(\boldsymbol{J}_{\boldsymbol{\theta}}^{(m+1)}\right)^{\top}, \tag{21a}$$

$$\boldsymbol{J}_{\boldsymbol{\theta}}^{(m+1)} = \frac{\partial \boldsymbol{Y}_{m+1}}{\partial \boldsymbol{\theta}} = \frac{\partial}{\partial \boldsymbol{\theta}}\Big[\text{NAO}(\text{TE}(\boldsymbol{Y}_{m-n+1:m}; \boldsymbol{\theta}_{\text{TE}}), \boldsymbol{X}, \boldsymbol{\theta}_{\text{NAO}})\Big], \tag{21b}$$

$$p(\boldsymbol{Y}_{m+1}|\mathcal{D}) \approx \mathcal{N}(\boldsymbol{Y}_{m+1}(\boldsymbol{\theta}_{\text{MAP}}), \boldsymbol{J}_{\boldsymbol{\theta}}^{(m+1)} \boldsymbol{\Sigma} (\boldsymbol{J}_{\boldsymbol{\theta}}^{(m+1)})^{\top}). \tag{21c}$$

These equations quantifies how epistemic uncertainty, represented by the parameter covariance $\boldsymbol{\Sigma}$, propagates through ASNO to the predicted temperature fields. Equation (21a) computes the predictive covariance by linearly mapping parameter uncertainty through the network Jacobian $\mathbf{J}_{\boldsymbol{\theta}}^{(m+1)}$, which measures the sensitivity of predictions to parameter changes. Equation (21c) then defines a Gaussian predictive distribution of the melt pool window. This provides epistemic uncertainty estimates of both mean and variance, enabling ASNO to quantify melt pool window prediction confidence and aids uncertainty-aware decisions in the BO framework.

3.3. Bayesian Optimization for Concurrent Design

Following the surrogate modeling and uncertainty quantification sections, this section presents the complete UNiCO-AM optimization framework. The framework initiates by training the uncertainty-aware ASNO surrogate on 500 high-fidelity DED simulations generated via Latin Hypercube Sampling (LHS) over the normalized process-material design space. Each sample combines eleven laser power profile parameters with material properties including density, solidus and liquidus temperatures, and temperature-dependent property vectors $\boldsymbol{C}_p(T)$ and $\boldsymbol{k}(T)$. The dataset used for training is described in Section 2.4. Once trained, ASNO functions as a fast, uncertainty-aware surrogate capable of predicting full-field melt-pool temperature windows across diverse process-material combinations. These predictive means and covariances form the foundation of the BO (Fig. 1), which integrates ASNO's predictions with physics-based evaluations from the *GAMMA* simulator to efficiently explore the joint design space.

ASNO predicts an initial batch of melt-pool temperature windows, each represented as a Gaussian distribution:

$$p\left(\boldsymbol{Y}_{m+1} \mid \mathcal{D}\right) \approx \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{Y}_{m+1}}, \boldsymbol{\Sigma}_{\boldsymbol{Y}_{m+1}}), \tag{22}$$

where $\boldsymbol{\mu}_{\boldsymbol{Y}_{m+1}}$ and $\boldsymbol{\Sigma}_{\boldsymbol{Y}_{m+1}}$ denote the predictive mean and epistemic uncertainty of the melt pool window, respectively. These outputs are used to compute key thermal metrics such as the normalized effective cooling rate $\widehat{T}_{\text{eff}}$, its coefficient of variation $CV_{\widehat{T}}$, and the physical feasibility constraints g_i as shown in Section 2.4.

These thermal metrics are passed to the acquisition function, which balances *exploration*, guided by ASNO's predictive uncertainty, and *exploitation*, driven by objective improvement. In UNiCO-AM the acquisition function selected is Expected Hypervolume Improvement (EHVI), chosen because it naturally extends single-objective improvement to multi-objective trade-off optimization and effectively guides the search toward Pareto-optimal process-material combinations. The two thermal objectives are the mean normalized cooling rate $\mu_{\widehat{T}}(\mathbf{x})$ to be maximized and the mean coefficient of variation $\mu_{CV}(\mathbf{x})$ to be minimized, as defined in Eq. (8). The Expected Hypervolume Improvement (EHVI) quantifies the expected expansion of the Pareto front across these objectives by measuring how much a new candidate design $\boldsymbol{x}$ improves the trade-off region between cooling intensity and thermal stability. The Expected Hypervolume Improvement (EHVI) is defined as:

$$\text{EHVI}(\mathbf{x}) = \mathbb{E}\Big[\mathcal{H}\Big(\mathcal{P} \cup \{\mu_{\widehat{T}}(\mathbf{x}), -\mu_{CV}(\mathbf{x})\}\Big) - \mathcal{H}(\mathcal{P})\Big], \tag{23}$$

where $\mathcal{P}$ denotes the current set of non-dominated (Pareto-optimal) objective values, and the Pareto hypervolume $\mathcal{H}(\mathcal{P})$ is computed as

$$\mathcal{H}(\mathcal{P}) = \text{Vol}\left(\bigcup_{[y_1,y_2]\in\mathcal{P}} [y_1, r_1] \times [y_2, r_2]\right), \tag{24}$$

with $[y_1, y_2] = [\mu_{\widehat{T}}(\mathbf{x}), -\mu_{CV}(\mathbf{x})]$ and $\boldsymbol{r} = [r_1, r_2]$ as the reference point. Maximizing EHVI($\mathbf{x}$) therefore drives the search toward process-material combinations that increase cooling efficiency while reducing thermal variability, ensuring thermally optimal and statistically robust designs. Compared to traditional GP-based Bayesian optimization or evolutionary search methods, the neural operator surrogate in UNiCO-AM scales more effectively to the high-dimensional joint process–material design space and leverages spatio-temporal field prediction to extract richer information per simulation, enabling practical sample-efficient exploration where kernel-based GP models and purely black-box optimizers become computationally and statistically limiting.

For batch optimization, the q-EHVI formulation extends Eq. (23) to evaluate multiple candidates $\{\mathbf{x}_j\}_{j=1}^{q}$ in parallel:

$$q\text{-EHVI} = \mathbb{E}\Big[\mathcal{H}\Big(\mathcal{P} \cup \{\mu_{\widehat{T}}(\mathbf{x}_j), -\mu_{CV}(\mathbf{x}_j)\}_{j=1}^{q}\Big) - \mathcal{H}(\mathcal{P})\Big], \tag{25}$$

where $q = 64$ denotes the batch size, corresponding to 64 distinct process-material combinations generated per iteration, each defined by a unique pairing of laser power profile parameters and temperature-dependent material properties.

Feasibility is incorporated probabilistically through chance constraints, which enforce that physical limits are satisfied with a specified confidence level.

$$w(\mathbf{x}) = \prod_{i=1}^{K} \Pr\left(g_i(\mathbf{x}) \le 0\right), \tag{26}$$

assigning a feasibility weight $w(\mathbf{x})$ to each candidate. The constrained Expected Hypervolume Improvement (cEHVI) [35] is then computed as

$$\text{cEHVI}(\mathbf{x}) = w(\mathbf{x})\,\text{EHVI}(\mathbf{x}), \tag{27}$$

ensuring that only physically valid and manufacturable combinations contribute to Pareto improvement.

To promote diversity and prevent redundant evaluations, local penalization is applied in the process-material design space:

$$\|\mathbf{x}_j - \mathbf{x}_k\|_2 > \rho_k, \qquad \forall\, j \neq k, \quad \mathbf{x}_j, \mathbf{x}_k \in \{\mathbf{x}_p, \mathbf{x}_m\}, \tag{28}$$

where $\mathbf{x}_j$ and $\mathbf{x}_k$ denote two distinct candidate vectors, each combining laser process parameters $\mathbf{x}_p$ and material properties $\mathbf{x}_m$. The minimum separation radius ρ_k is derived from the Lipschitz continuity of the thermal objectives $\mu_{\widehat{T}}(\mathbf{x})$ and $\mu_{CV}(\mathbf{x})$, ensuring that each batch samples physically distinct and informative regions of the design space. Each candidate batch is then evaluated in parallel using the *GAMMA* simulator to obtain new full-field thermal responses $\boldsymbol{Y}$, which are subsequently used to update ASNO's predictive mean and Laplace-based covariance estimates.

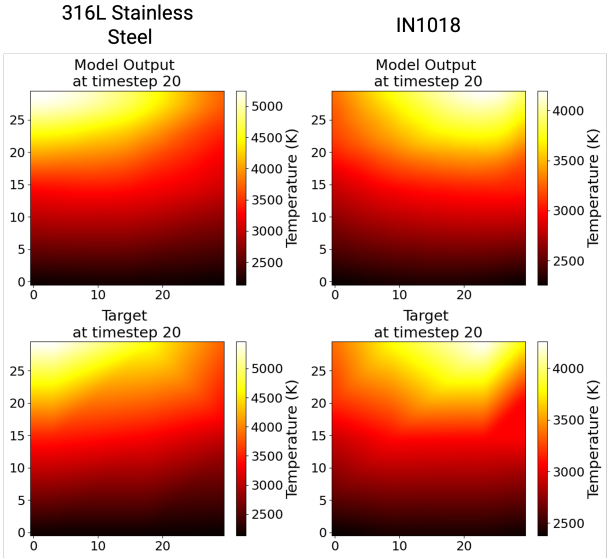


Figure 4. ASNO model melt pool window predictions (model output) align with GAMMA outputs (targets) across time.

After each batch evaluation, ASNO is fine-tuned using the Low-Rank Adaptation (LoRA) method, enabling efficient parameter updates without catastrophic forgetting. All design variables are normalized to the range $[0, 1]$, and optimization proceeds until the relative hypervolume improvement $\Delta\mathcal{H}/\mathcal{H}_{\text{prev}}$ falls below 0.1% or until a maximum of 400 iterations is reached. This iterative feedback loop, consisting of prediction, acquisition, simulation, and model updates, achieves process and material designs that are thermally optimized, stable, and physically feasible.

4. Results and Discussion

In this section, UNiCO-AM is evaluated along two test cases: (i) the predictive accuracy and cross-alloy generalization of the ASNO surrogate, and (ii) the results of the concurrent BO in optimizing high-cooling, low-variability, and feasible designs. We test ASNO's ability to reconstruct spatio-temporal thermal fields for both seen and unseen alloys, and assess the end-to-end BO process that uses these predictions to propose optimized process-material combinations.

4.1. Surrogate Accuracy and Cross-Material Generalization

ASNO predictions are evaluated along with the true predictions from GAMMA for two training alloys (316L stainless steel and 1018 steel) and one held-out alloy (17-4 PH stainless steel). For the in-distribution alloys, Figure 4 shows strong similarity between ASNO predictions and ground-truth temperature fields. The temporal extrapolator captures heating-cooling progression, while the nonlocal attention operator restores sharp spatial gradients, resulting accurate reconstruction of the melt-pool window. ASNO achieves a mean absolute percentage error (MAPE) of 4.5% and a test loss of 0.0140, confirming its high-fidelity reconstruction of transient melt-pool dynamics and accurate alignment with physics-based thermal behavior observed in GAMMA simulations. To further validate ASNO's performance, Table 1 presents a comparative evaluation against other state-of-the-art neural operator architectures, including FNO, DeepONet, UNet, and GNOT [33]. ASNO not only achieves the lowest test loss among all models but also maintains a favorable balance between model complexity and training efficiency. In addition to loss values, we report the Structural Similarity Index Measure (SSIM) as a field-level structural fidelity metric that quantitatively reflects spatial agreement of predicted melt-pool temperature patterns. This demonstrates that ASNO provides better predictive accuracy and computational efficiency, making it a highly effective surrogate model for additive manufacturing applications.

Table 1. Comparison of models on the Additive Manufacturing dataset.

Model Type	Trainable Params	Time per Epoch (s)	Best Test Loss	SSIM ↑
ASNO	2,825,444	6.17	**0.0140**	**0.93**
FNO	3,654,312	7.10	0.0357	0.86
DeepONet	3,385,605	5.25	0.0456	0.82
UNet	4,564,157	6.64	0.0652	0.75
GNOT	**2,442,678**	6.60	0.0332	0.88
Transolver	2,636,9158	8.71	0.0502	0.80
Linear + NAO	2,943,435	4.75	0.0397	0.84
Transformer Encoder (TE)	2,784,920	**3.95**	0.0563	0.77

One important aspect of ASNO is its ability to capture the underlying physics of the DED process, which can be validated through its generalizability. Generalizability refers to the model's ability to predict meltpool temperature window for unseen materials and process conditions, reflecting true physical understanding rather than data memorization. Here, we assess this by testing ASNO on the unseen 17-4 PH stainless steel alloy, which differs significantly in thermophysical properties from the training alloys, enabling evaluation of its extrapolation capability across materials. Figure 5 shows that 17-4 PH occupies a distinct region in the latent material property space, clearly separated from the stainless steel (316L, 1018) and nickel-based superalloy (IN625, IN718) clusters used during training. The latent space is constructed by applying Principal Component Analysis (PCA) to the 17 material variables that define one material and reducing them to two dimensions, effectively illustrating the distinct thermophysical nature of 17-4 PH. This separation arises from its higher solidus and liquidus temperatures, greater density, and moderate thermal conductivity, which together produce steeper thermal gradients and faster cooling transients. When evaluated without fine-tuning (Zero-Shot evaluation), ASNO predicts coherent meltpool window, accurately capturing hot spots, melt-pool contraction, and global heat dissipation patterns in the melt pool window. The out-of-distribution MAPE of 19.5% and test loss of 0.0356 remain within acceptable limits for BO-guided exploration. These results demonstrate that ASNO effectively leverages the joint process-material embedding, which aids exploration and exploitation in the UNiCO-AM framework.

4.2. Optimization Performance

Having discussed in the previous section that ASNO reconstructs melt-pool temperature windows for unknown materials efficiently, we now examine UNiCO-AM's capabilities for concurrent material–process optimization. The UNiCO-AM utilizes the ASNO predictive means and variances of the melt-pool

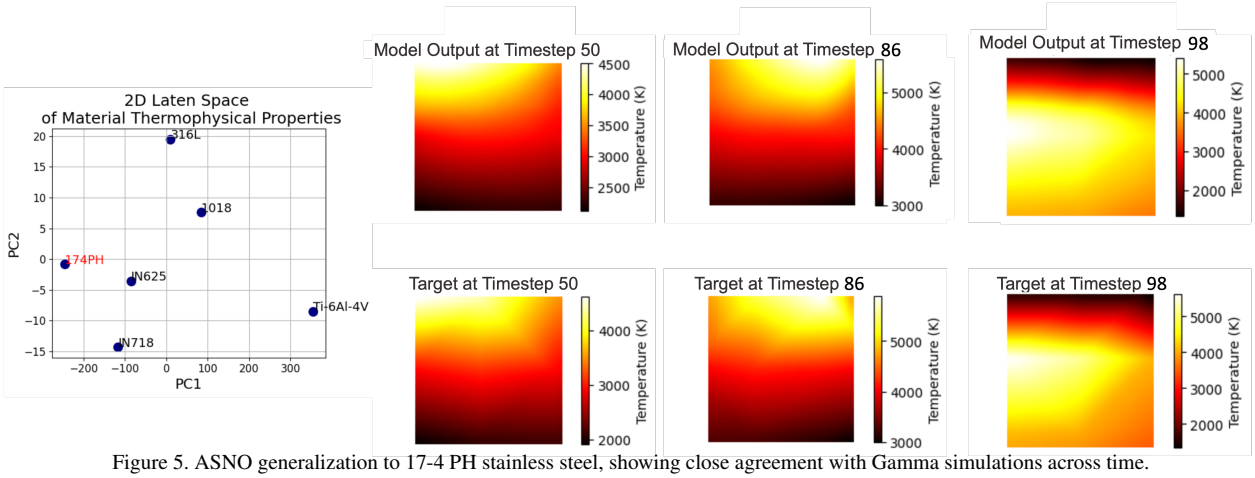


Figure 5. ASNO generalization to 17-4 PH stainless steel, showing close agreement with Gamma simulations across time.

temperature window at each rollout step and converts them into objective metrics, specifically the normalized effective cooling rate $\dot{T}_{\text{eff}}$ and its coefficient of variation $\text{CV}_{\dot{T}}$, and selects candidate designs whose expected improvement expands the current Pareto set while following feasibility probabilities. Because the q-EHVI acquisition explicitly depends on both predictive mean and epistemic variance, regions exhibiting elevated structural prediction error are associated with increased uncertainty estimates from the LLA posterior, which promotes exploration rather than misleading exploitation, a behavior further supported by the added SSIM-based structural fidelity evaluation in Table 1.

Figure 6 illustrates the Pareto front obtained from the UNiCO-AM concurrent optimization loop, which jointly maximizes the effective normalized solidus cooling rate and minimizes its coefficient of variation. The resulting front reveals two dominant thermal performance regimes within the studied design space: Inconel-based candidates occupy the left region, exhibiting lower cooling rates but reduced coefficient of variation due to higher heat retention and slower thermal diffusion, whereas stainless steel candidates lie on the right, achieving faster cooling at the expense of increased variability owing to their higher thermal conductivity and greater sensitivity to laser power fluctuations. We note that this Pareto clustering reflects the representative alloys studied here, consistent with known differences in alloy-class thermal diffusivity and cooling behavior in additive manufacturing [36], rather than a universal material-family generalization. The smooth progression along the Pareto front indicates that the optimization effectively explores the coupled process–material space, transitioning from stable, diffusion-controlled cooling modes to more aggressive, conductivity-driven regimes. Candidate selection is weighted by constraint satisfaction probabilities derived from Laplace-based uncertainty estimates, ensuring exploration remains physically feasible while still uncovering high-performance regions. The BO loop executed 400 iterations, each evaluating 64 process–material candidates with the high-fidelity *GAMMA* simulator on four GPUs, resulting in approximately 25,600 simulations that progressively refined the surrogate and Pareto front. Importantly, the optimized thermophysical trends along the Pareto front should be interpreted as

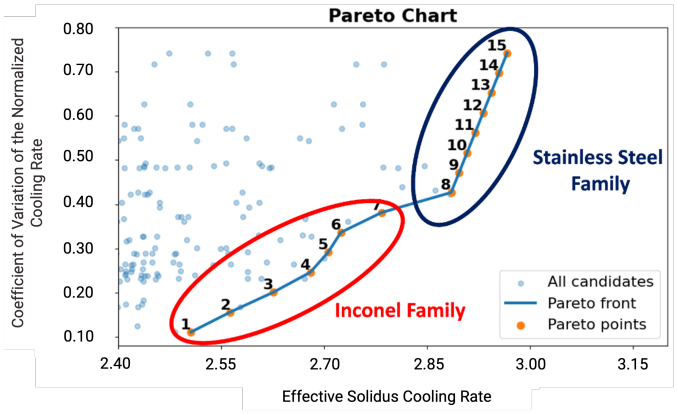


Figure 6. Pareto front illustrating the trade-off between solidus cooling rate and variability and variability in cooling rate

inverse thermal-performance targets that guide alloy screening and process-aware material selection, rather than as directly prescribable material-property curves. In practice, these results can assist materials scientists by identifying desirable cooling-rate regimes and mapping them to the closest realizable alloy classes or by informing directional alloy or process development strategies that achieve comparable thermal responses. Overall, the Pareto front provides actionable design insights, highlighting stainless steel configurations as favorable for achieving rapid solidification, enhanced grain refinement, and high thermal responsiveness, while Inconel-based designs remain better suited for stability-driven, defect-sensitive builds in metal additive manufacturing.

As the material properties are defined in a continuous parameter space, *UNiCO-AM* is capable of generating novel material property combinations that were not explicitly present in the training dataset, enabling discovery beyond discrete alloy compositions. To further analyze the representative co-design Pareto front, we focus on Pareto points 8 and 9, which lie within the stainless steel cluster and represent a balance between accelerated cooling and acceptable variance. These designs perform better compared to their Inconel counterparts in terms of solidification intensity while avoiding the extreme fluctuations seen in more aggressive laser strategies. Figure 7 shows the corresponding laser power profiles, where Pareto point 8 exhibits a smoother, ramp-based modulation, whereas point 9 employs

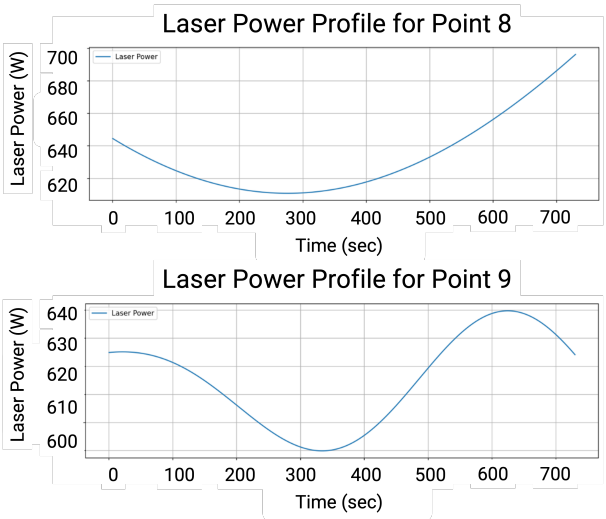


Figure 7. Laser power profiles corresponding to Pareto points 8 (top) and 9 (bottom).

Table 2. Bulk material properties for Pareto points 8 and 9.

Point	Density (g/cm^3)	Solidus (K)	Liquidus (K)
8	7.80	1673	1713
9	7.99	1648	1673

mildly oscillatory adjustments that induce larger thermal fluctuations. The initial power reduction phase observed in both profiles mitigates excessive melt-pool expansion during early heating, while their late-stage increases compensate for rising conductive losses as thermal gradients equilibrate. This behavioral contrast is consistent with the uncertainty trends reported earlier: point 9's stronger oscillations align with slightly higher predictive variance around local peaks of the cooling curve, whereas point 8 remains closer to the central credible band.

The material properties corresponding to points 8 and 9, shown in Table 2 and Figure 8, reveal that subtle differences in heat capacity and thermal conductivity over the solidification interval significantly influence the cooling curve shape and coefficient of variance magnitude. Point 9 exhibits slightly higher conductivity fluctuations at elevated temperatures, which correlates with the observed increase in variability relative to point 8.

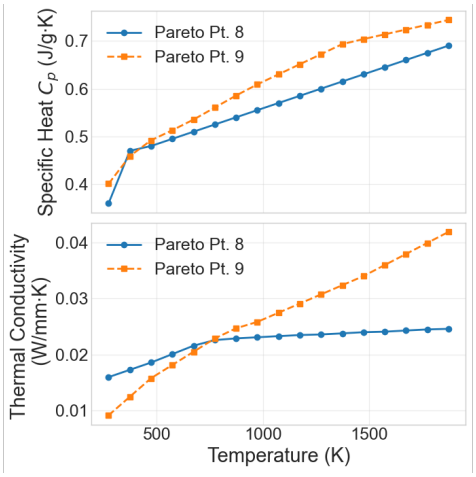


Figure 8. Temperature-dependent specific heat and thermal conductivity for Pareto points 8 and 9, illustrating smoother variation for point 8 and steeper high-temperature gradients for point 9.

This correlation arises because variations in high-temperature conductivity alter local heat extraction efficiency, leading to uneven cooling across the melt pool and thereby amplifying the observed cooling-rate variance. These results confirm that the UNiCO-AM balances material process variables, identifying profiles that exploit both property-driven and control-driven thermal regulation mechanisms. Because the UNiCO-AM loop is continuously re-informed by ASNO's updated MAP parameters, the search remains statistically grounded as it moves across alloys and laser power profile families.

5. Conclusion

In this paper, we developed *UNiCO-AM*, an uncertainty-aware neural co-design framework that jointly designs the material and process conditions in metal additive manufacturing. The primary objective of the UNiCO-AM framework is to discover new material-process combinations in additive manufacturing that lead to improved material properties in the fabricated part. The paper presents three main contributions: it introduces a physics-informed ASNO capable of accurately modeling full-field thermal dynamics across diverse alloys and process conditions; it integrates Laplace-based uncertainty quantification to provide epistemic uncertainty estimates for confidence-aware predictions; and it develops a constrained multi-objective Bayesian optimization framework in a high-dimensional design space that efficiently identifies Pareto-optimal process-material combinations balancing cooling efficiency, thermal stability, and manufacturability.

First, we developed an ASNO architecture specifically designed for material-process modeling in additive manufacturing. ASNO demonstrated the ability to learn the underlying physics governing the additive manufacturing process, achieving the MAPE of 4.5%, outperforming conventional neural operator methods. Furthermore, it exhibited strong out-of-distribution generalization when tested on 17-4 PH stainless steel, achieving an error of 19.5%, thereby showcasing its capability to model unseen material-process combinations. The improved performance of ASNO arises from its ability to capture the coupled spatio-temporal thermodynamic behavior intrinsic to additive manufacturing, which is critical for accurately predicting temperature evolution and microstructural dynamics. Additionally, this work is the first to demonstrate that epistemic uncertainty can be quantified for neural operators, enabling their use as uncertainty-aware surrogates within Bayesian optimization frameworks.

Second, we completed the *UNiCO-AM* framework by developing the final component of the Bayesian optimization, introducing a q-EHVI acquisition function that enables concurrent material-process optimization in additive manufacturing. The *UNiCO-AM* framework successfully achieves distinct Pareto trade-offs between cooling rate and cooling variance, identifying material and process-dependent thermal behaviors that enhance the microstructure and hardness of the fabricated parts. Notably, we observed that stainless steel alloys exhibit more favorable performance compared to their Inconel counterparts, demonstrating the effectiveness of material-process co-design optimization within *UNiCO-AM*. Overall, this work establishes a generalized methodology capable of performing concurrent

material-process co-design optimization for a wide range of tasks in additive manufacturing by changing the objective function. Future work will focus on extending the material domain from metals to polymers and other classes of materials, while also enabling concurrent geometry-material-process optimization. To achieve geometry-driven co-design, the ASNO architecture will be further enhanced to learn and generalize across varying geometric configurations, allowing the framework to capture geometry-dependent thermal and microstructural behaviors in additive manufacturing.

Acknowledgements

We appreciate the grant support from the NSF HAMMER-ERC (Engineering Research Center for Hybrid Autonomous Manufacturing, Moving from Evolution to Revolution) under Award EEC-2133630 and the NSF MADE-PUBLIC Future Manufacturing Research Grant Program under Award CMMI-2037026, NSF Boosting Research Ideas for Transformative. Yi-Ping Chen appreciates the Taiwan–Northwestern Doctoral Scholarship funded by the Ministry of Education in Taiwan.

References

[1] W. L. Ng, G. L. Goh, G. D. Goh, J. S. J. Ten, W. Y. Yeong, Progress and opportunities for machine learning in materials and processes of additive manufacturing, Advanced Materials 36 (34) (2024) 2310006.

[2] V. Karkaria, Y.-K. Tsai, Y.-P. Chen, W. Chen, An optimization-centric review on integrating artificial intelligence and digital twin technologies in manufacturing, Engineering Optimization 57 (1) (2025) 161–207.

[3] T. DebRoy, T. Mukherjee, H. Wei, J. Elmer, J. Milewski, Metallurgy, mechanistic models and machine learning in metal printing, Nature Reviews Materials 6 (1) (2021) 48–68.

[4] V. Karkaria, A. Goeckner, R. Zha, J. Chen, J. Zhang, Q. Zhu, J. Cao, R. X. Gao, W. Chen, Towards a digital twin framework in additive manufacturing: Machine learning and bayesian optimization for time series process optimization, Journal of Manufacturing Systems 75 (2024) 322–332.

[5] Y.-P. Chen, V. Karkaria, Y.-K. Tsai, F. Rolark, D. Quispe, R. X. Gao, J. Cao, W. Chen, Real-time decision-making for digital twin in additive manufacturing with model predictive control using time-series deep neural networks, Journal of Manufacturing Systems 80 (2025) 412–424.

[6] Y.-P. Chen, Y.-K. Tsai, V. Karkaria, W. Chen, Uncertainty-aware digital twins: Robust model predictive control using time-series deep quantile learning, Journal of Mechanical Design 148 (2) (2026) 021702.

[7] S. Yang, Y. F. Zhao, Additive manufacturing-enabled design theory and methodology: a critical review, The International Journal of Advanced Manufacturing Technology 80 (1) (2015) 327–342.

[8] G. Liu, Y. Xiong, D. W. Rosen, Multidisciplinary design optimization in design for additive manufacturing, Journal of Computational Design and Engineering 9 (1) (2022) 128–143.

[9] E. Ulu, R. Huang, L. B. Kara, K. S. Whitefoot, Concurrent structure and process optimization for minimum cost metal additive manufacturing, Journal of Mechanical Design 141 (6) (2019) 061701.

[10] V. Mishra, C. Ayas, M. Langelaar, Design for material properties of additively manufactured metals using topology optimization, Materials & Design 235 (2023) 112388.

[11] R. Ranjan, Z. Chen, C. Ayas, M. Langelaar, F. Van Keulen, Overheating control in additive manufacturing using a 3d topology optimization method and experimental validation, Additive Manufacturing 61 (2023) 103339.

[12] S. Liao, S. Webster, D. Huang, R. Council, K. Ehmann, J. Cao, Simulation-guided variable laser power design for melt pool depth control in directed energy deposition, Additive Manufacturing 56 (2022) 102912.

[13] M. Mozaffar, S. Liao, J. Jeong, T. Xue, J. Cao, Differentiable simulation for material thermal response design in additive manufacturing processes, Additive Manufacturing 61 (2023) 103337.

[14] H. Chen, X. Liu, X. Liu, P. Witherell, Meltpoolgan: Melt pool prediction from path-level thermal history, Additive Manufacturing 84 (2024) 104095.

[15] A. Kao, T. Gan, C. Tonry, I. Krastins, K. Pericleous, Thermoelectric magnetohydrodynamic control of melt pool dynamics and microstructure evolution in additive manufacturing, Philosophical Transactions of the Royal Society A 378 (2171) (2020) 20190249.

[16] H. Liu, Y.-S. Ong, X. Shen, J. Cai, When gaussian process meets big data: A review of scalable gps, IEEE transactions on neural networks and learning systems 31 (11) (2020) 4405–4423.

[17] D. Patel, R. Sharma, Y. Guo, Computational, data-driven, and physics-informed machine learning approaches for microstructure modeling in metal additive manufacturing, arXiv preprint arXiv:2505.01424.

[18] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, A. Anandkumar, Fourier neural operator for parametric partial differential equations, arXiv preprint arXiv:2010.08895.

[19] I. Y. Kim, O. L. de Weck, Adaptive weighted sum method for multiobjective optimization: a new method for pareto front generation, Structural and multidisciplinary optimization 31 (2) (2006) 105–116.

[20] Z. Deng, F. Zhou, J. Zhu, Accelerated linearized laplace approximation for bayesian deep learning, Advances in Neural Information Processing Systems 35 (2022) 2695–2708.

[21] K. S. Ramani, C. He, Y.-L. Tsai, C. E. Okwudire, E. Malekipour, Smartscan: An intelligent scanning approach for uniform thermal distribution, reduced residual stresses and deformations in pbf additive manufacturing, Additive Manufacturing 52 (2022) 102643.

[22] H. Hmede, S. Nicaise, L. Paquet, A. Wehbe, Optimal control of the temperature by the laser path and the thermal treatment time in selective laser melting process (2023).

[23] T. Mukherjee, W. Zhang, T. DebRoy, An improved prediction of residual stresses and distortion in additive manufacturing, Computational Materials Science 126 (2017) 360–372.

[24] D. Rosenthal, The theory of moving sources of heat and its application to metal treatments, Transactions of the American Society of Mechanical Engineers 68 (8) (1946) 849–865.

[25] T. DebRoy, H. L. Wei, J. S. Zuback, T. Mukherjee, J. W. Elmer, J. O. Milewski, A. M. Beese, A. Wilson-Heid, A. De, W. Zhang, Additive manufacturing of metallic components—process, structure and properties, Progress in Materials Science 92 (2018) 112–224. doi:10.1016/j.pmatsci.2017.10.001.

[26] S. A. Khairallah, A. T. Anderson, A. Rubenchik, W. E. King, Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones, Acta Materialia 108 (2016) 36–45.

[27] J. Bennett, J. Glerum, J. Cao, Relating additively manufactured part tensile properties to thermal metrics, CIRP Annals 70 (1) (2021) 187–190.

[28] O. A. Ezekoye, Conduction of heat in solids, Springer, 2016.

[29] L. Meyer, G. Witt, Investigation of the track width-dependent melt pool characteristics during laser-sintering of polyamide 12 in correlation to various focus diameters, Progress in Additive Manufacturing 5 (1) (2020) 19–25.

[30] N. Baldi, A. Giorgetti, M. Palladino, I. Giovannetti, G. Arcidiacono, P. Citti, Study on the effect of preheating temperatures on melt pool stability in inconel 718 components processed by laser powder bed fusion, Metals 13 (10) (2023) 1792.

[31] Y.-K. Tsai, R. J. Malak, Jr, Control co-design with performance-robustness trade-off using tube-based stochastic model predictive control, ASME Letters in Dynamic Systems and Control 5 (3) (2025) 030903.

[32] Y.-K. Tsai, R. J. Malak Jr, Surrogate-assisted constraint-handling technique for parametric multi-objective optimization, Structural and Multidisciplinary Optimization 67 (9) (2024) 167.

[33] V. Karkaria, D. Lee, Y.-P. Chen, Y. Yu, W. Chen, An attention-based spatio-temporal neural operator for evolving physics, Machine Learning: Science and Technology.

[34] A. H. Committee, et al., Properties and selection: nonferrous alloys and special-purpose materials, ASM international, 1990.

[35] M. Abdolshah, A. Shilton, S. Rana, S. Gupta, S. Venkatesh, Expected hypervolume improvement with constraints, in: 2018 24th International Conference on Pattern Recognition (ICPR), IEEE, 2018, pp. 3238–3243.

[36] T. Mukherjee, V. Manvatkar, A. De, T. DebRoy, Dimensionless numbers in additive manufacturing, Journal of Applied Physics 121 (6).