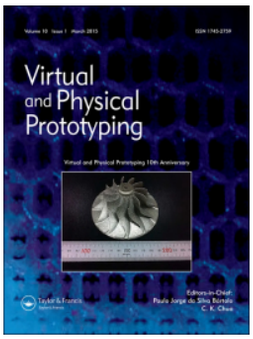

Virtual and Physical Prototyping

ISSN: 1745-2759 (Print) 1745-2767 (Online) Journal homepage: www.tandfonline.com/journals/nvpp20

Heat treatment design for additively manufactured nickel superalloy via Bayesian optimization of ultimate tensile strength and energy consumption

Enjy Katary, Kinga Sekuła, Dawid Krok, Marcin Posmyk, Michał Dusza, Wojciech Matusik & Konrad Gruber

To cite this article: Enjy Katary, Kinga Sekuła, Dawid Krok, Marcin Posmyk, Michał Dusza, Wojciech Matusik & Konrad Gruber (2026) Heat treatment design for additively manufactured nickel superalloy via Bayesian optimization of ultimate tensile strength and energy consumption, Virtual and Physical Prototyping, 21:1, e2686064, DOI: 10.1080/17452759.2026.2686064

To link to this article: https://doi.org/10.1080/17452759.2026.2686064

VIRTUAL AND PHYSICAL PROTOTYPING
2026, VOL. 21, NO. 1, e2686064 (25 pages)
https://doi.org/10.1080/17452759.2026.2686064

Heat treatment design for additively manufactured nickel superalloy via Bayesian optimization of ultimate tensile strength and energy consumption

Enjy Katary [a,b], Kinga Sekuła [a], Dawid Krok [a], Marcin Posmyk [a], Michał Dusza [a], Wojciech Matusik [c] and Konrad Gruber [a]

[a]Department of Advanced Manufacturing Technologies, Faculty of Mechanical Engineering, Wroclaw University of Science and Technology, Wroclaw, Poland; [b]University of Lille, CNRS, Centrale Lille, UMR 9013 - LaMcube - Laboratoire de mécanique multiphysique et multiéchelle, Lille, France; [c]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

Heat treatment is essential for meeting mechanical-performance requirements in additively manufactured components, but it consumes significant furnace energy and involves coupled time–temperature interactions that make trial-and-error development inefficient. We demonstrate a multi-objective Bayesian optimization (MOBO) protocol for Laser Powder Bed Fused (LPBF) Inconel 718 (IN718) that simultaneously maximises ultimate tensile strength (UTS) and minimises estimated cycle energy consumption. Profilometry-Based Indentation Plastometry (PIP) provides rapid UTS estimates during the campaign, while a calibrated furnace model estimates per-cycle energy consumption (EC). Relative to the average seed baseline, the campaign achieved a 21% increase in UTS and a 118% improvement in the energy efficiency objective (1/EC), with a Pareto hit ratio of 0.33. The identified schedules favour short solution annealing (15 min, 950–1034°C) paired with moderate aging, and microstructural observations are consistent with the inferred trends, indicating that heat treatment primarily modifies segregation products and precipitate populations rather than grain morphology. Within practical processing constraints, the proposed workflow makes resource intensity an explicit optimisation target while retaining metallurgical interpretability and practical experimental cadence.

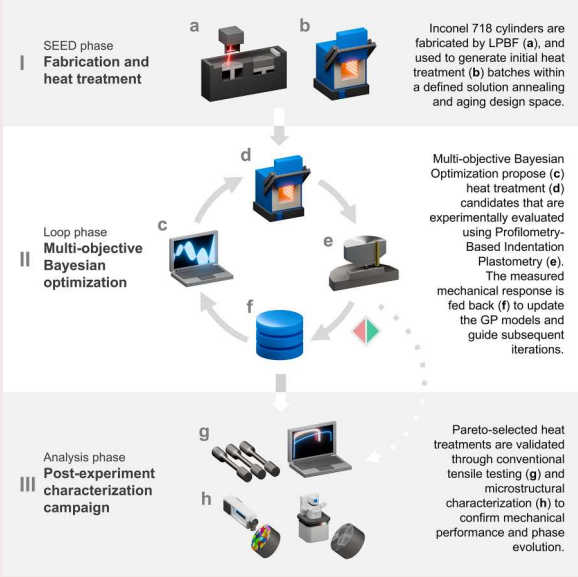

ARTICLE HISTORY
Received 2 March 2026
Accepted 11 May 2026

KEYWORDS
Multi-objective Bayesian optimization; profilometry-based indentation plastometry; heat treatment; machine learning; laser powder bed fusion

CONTACT Konrad Gruber konrad.gruber@pwr.edu.pl Department of Advanced Manufacturing Technologies, Faculty of Mechanical Engineering, Wroclaw University of Science and Technology, Lukasiewicza 5, 50-371 Wroclaw, Poland

Supplemental data for this article can be accessed online at https://doi.org/10.1080/17452759.2026.2686064.

1. Introduction

Additive manufacturing (AM) enables complex metallic components with shortened supply chains and rapid iteration compared with conventional routes. Among metal AM, Laser Powder Bed Fusion (LPBF) is the most widespread AM modality for high-value nickel superalloys because it offers fine feature resolution and near-net shape capability, although with steep thermal gradients and rapid solidification that imprint microstructural heterogeneity [1].

Inconel 718 (IN718) produced by LPBF typically solidifies into a cellular/dendritic γ matrix decorated by segregation-induced phases; the as-built state also carries high residual stresses from cyclic heating and cooling inherent to the layerwise process [2]. Post-build heat treatment is therefore required to relieve residual stresses, reduce microsegregation, and re-establish a phase balance suitable for subsequent strengthening. The key levers are the dissolution of Laves and δ phases and the controlled precipitation of γ'/γ'' during aging, which together determine the achievable yield strength, ultimate tensile strength, and ductility [3].

A wide range of post-processing heat-treatment schedules has been explored for additively manufactured Inconel 718, particularly for LPBF-produced material [4, 5]. In practice, however, most studies evaluate only a small set of conditions, typically two to five distinct cycles, because furnace access, sample throughput, and mechanical testing time are limiting factors. For example, Fayed et al. examined five combinations of homogenisation time (1, 4, 7 h at 1080°C) and solution time (15–60 min at 980°C) to assess the high-temperature tensile response [6]. Salgado-López et al. varied the first aging hold between 6, 8, and 10 h at 720°C and reported a hardness maximum near 8 h (∼540 HV), with overaging and a drop to ∼516 HV at 10 h [7].

Other work has targeted cycle simplification. Benzing et al. proposed an abbreviated route for LPBF IN718 that consolidates multiple steps into a single high-pressure, high-temperature stage (HIP at 1020°C), followed by single aging at 720°C [8]. This reduced the total heat-treatment time from roughly 42 h to 15 h while maintaining mechanical performance. Dwivedi and Khurana applied a response-surface methodology to optimise homogenisation time (1–8 h at 1080°C) and solution temperature (980–1140°C for 1 h), followed by fixed standard aging [9]. Their statistical model identified an intermediate regime (about 2 h at 1080°C with a 980°C solution) as a practical compromise between competing responses.

These studies are limited to exploring only a small portion of the time–temperature space, with trade-offs between properties, processing time, and energy use rarely examined directly. Furnace power demand links heat-treatment decisions to cost, throughput, and environmental burden. At the same time, higher temperatures or longer holds do not guarantee superior strength, since dissolution, precipitation, and overaging kinetics introduce non-monotonic responses. The industrial relevance of this problem is reinforced by the broader AM process chain: post-processing accounts for a substantial fraction of metal AM production cost, with approximately 27% attributed to operations such as support removal, surface finishing, coating, and heat treatment [10]. Because the necessity and extent of these steps depend strongly on the alloy, AM route, and application, unnecessary or overly conservative post-processing can erode the economic benefit of AM. This makes heat-treatment design not only a metallurgical problem, but also a process-chain optimisation problem in which performance must be balanced against resource use, cost, and qualification effort.

Profilometry-Based Indentation Plastometry (PIP) offers rapid, small-volume estimation of tensile-equivalent stress–strain curves by inverting residual indent profiles using a calibrated plasticity model [11]. It is well-suited to AM materials because it samples localised microstructural states, enables directional mapping across build orientations, and can reveal anisotropy or porosity when probed in different regions [12,13]. A practical advantage is the minimal specimen size requirement: the indented region only needs to exceed the spherical indenter radius by a small margin [14], allowing tests on small coupons or directly on standard metallographic mounts. In this study, PIP is used mainly as a ranking tool to screen heat-treatment candidates before confirmatory tensile tests, reducing experimental time and material consumption while maintaining decision quality.

Bayesian optimization (BO) has increasingly been adopted in manufacturing and materials research as a practical strategy to navigate expensive and potentially noisy design spaces with limited experiments [15,16]. BO proceeds sequentially by fitting a probabilistic surrogate model to the available observations and using this model to propose the next experiment in a data-efficient manner [17,18]. Gaussian process (GP) surrogates are commonly used because they can represent nonlinear response surfaces while providing both a mean prediction and an uncertainty estimate that reflects data sparsity [19]. This uncertainty is exploited by acquisition functions that balance exploration and exploitation; in this work, we employ a GP Upper Confidence Bound (GP-UCB) to prioritise candidates that are predicted either to perform well or to provide information in

regions of high uncertainty [18,20,21]. In multi-objective settings, BO extends naturally by modelling each objective under uncertainty and iteratively refining the Pareto front of non-dominated trade-offs between competing metrics [22].

In additive manufacturing, BO has been applied to tune LPBF process windows toward defect-free fabrication and improved part quality, including optimisation of melt-pool parameters and defect density in alloys such as AZ31 [23], AA2024 [24], and IN718 [25]. BO has also been used at the build-planning stage to optimise part orientation and support-volume trade-offs [26]. Related multi-objective frameworks have been integrated with semi-autonomous AM platforms for formulation discovery, where algorithms propose feedstock mixtures and rapidly identify Pareto-optimal materials within a few tens of iterations [27]. Beyond part fabrication, BO has been applied to upstream powder manufacturing; for example, in gas atomisation of Ni-Co superalloy powders, melt temperature and gas pressure were tuned to increase fine-powder yield [28]. Pareto-active learning and BO-based approaches have also been used to jointly optimise LPBF processing and heat-treatment parameters for Ti-6Al-4V, efficiently exploring large candidate spaces and identifying conditions that improve the strength–ductility trade-off [29].

We therefore pose heat-treatment design as an explicit bi-criterion problem: maximising UTS_{PIP} while minimising estimated cycle energy consumption (EC). This formulation ties the optimisation to operationally relevant quantities, aligns with practical processing constraints, and reflects the competition between performance and resource intensity in LPBF IN718. The present work integrates multi-objective Bayesian optimisation with Profilometry-Based Indentation Plastometry as a rapid, small-volume mechanical proxy and a calibrated semi-empirical furnace energy model to explore a two-stage schedule (solution annealing and aging). Heat treatment is parameterised by four decision variables (two temperatures and two hold times). Gaussian process surrogates are trained to predict UTS_{PIP} and E_C, and experiments are proposed sequentially to trace attainable Pareto trade-offs. In contrast to more common AM heat-treatment optimisation studies that focus primarily on performance metrics [10], the present framework makes estimated energy consumption an explicit co-objective and thereby targets an industrially relevant trade-off between strength and process resource intensity. The main contributions are therefore an energy-aware Pareto formulation for AM post-processing, the integration of PIP into the MOBO loop enabling same-day experiment-to-model feedback, the identification and tensile validation of Pareto-optimal schedules that reduce estimated energy input by up to 60% relative to the most intensive seed schedules while achieving tensile UTS in the 1330–1340 MPa range, and the release of code and data in a supplementary repository to support transfer to other alloys and furnace configurations with appropriate recalibration of the proxy and furnace model.

2. Experiment design

2.1. Experimental workflow and optimisation strategy

The experimental design and strategy are outlined in Figure 1. The study workflow comprised three stages: (i) fabrication of Inconel 718 samples by Laser Powder Bed Fusion and generation of seed heat-treatment batches; (ii) iterative multi-objective Bayesian optimization (MOBO) loops integrating Profilometry-Based Indentation Plastometry, heat treatment, and candidate selection; and (iii) a post-MOBO validation campaign involving microstructural investigation and tensile testing of optimal solutions (Figure 1(a)).

LPBF samples were built as vertical cylinders (∅ 16 mm × 100 mm, Figure 1(b)) using a low-heat-input parameter set, a fine layer thickness, and a stripe-scan strategy with 67° rotation between layers (Table S1) [1]. This combination reduces thermal distortion and improves microstructural homogeneity relative to more conventional scan strategies. The build direction was aligned with the z-axis, and the cylinders were sectioned into ~10 mm slices perpendicular to the build direction for subsequent heat treatment and PIP testing. Gas-atomized Inconel 718 powder with a 10–45 µm particle-size distribution and predominantly spherical morphology was used for fabrication; a representative powder image is shown in Figure 1(c), and the chemical composition is provided in Table S2.

Heat-treatment schedules were explored within a bounded design space (Figure 1(d)). Solution annealing (SA) was varied between 950 and 1200°C for 15 min to 2 h, and artificial aging (AA) was varied between 600 and 860°C for 2–8 h. IN718 was selected as a model system because its LPBF microstructure and phase-transformation pathways are well documented [2,5]. This provides a clear basis for interpreting trends observed during optimisation and serves as a reference for extending the approach to other precipitation-strengthened superalloys. For new alloy chemistries, comparable bounds could be established using CALPHAD-based phase-stability predictions [30]. Here, limits were chosen based on the literature, to ensure

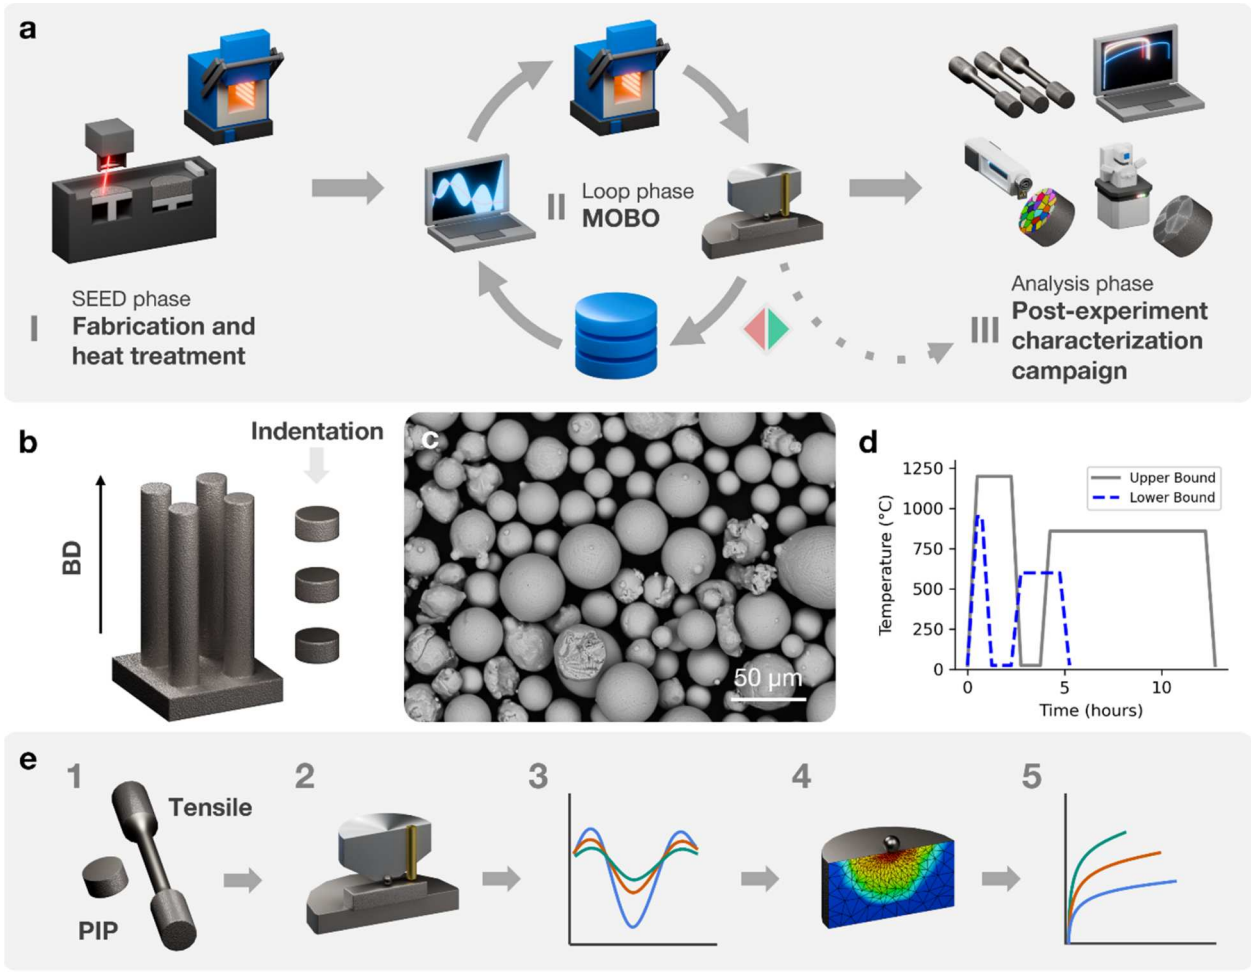

Figure 1. Experiment design and workflow. (a) Overall study framework comprising three main stages: (i) LPBF fabrication and seed batch generation,: (ii) iterative multi-objective Bayesian optimization (MOBO) integrating Profilometry-Based Indentation Plastometry (PIP), and (iii) post-MOBO validation campaign. (b) LPBF build configuration and sectioning scheme used for heat treatment and PIP testing. (c) Morphology of the IN718 powder used for sample fabrication. (d) Defined design space for solution annealing (SA) and artificial aging (AA) parameters explored during optimisation. (e) Schematic of the PIP method for rapid mechanical property estimation: (1) sample geometry comparison between PIP and tensile testing, (2) indentation, (3) profile measurement, (4) inverse FEM modelling, and (5) stress–strain curve reconstruction.

that the entire landscape of relevant transformations could be sampled: dissolution of Laves and δ/ σ phases during solution treatment and controlled precipitation of γ′/γ′′ during aging. These bounds, therefore, expose both high-temperature dissolution regimes and low-temperature precipitation regimes, enabling the optimiser to probe schedules that either suppress or promote the formation of typical IN718 phases.

The MOBO loop couples microstructural and mechanical testing with model-guided candidate generation. At each iteration, the Bayesian optimisation algorithm proposes new heat-treatment conditions, which are experimentally evaluated and then fed back into the models. Mechanical performance is assessed rapidly using PIP, where a spherical indenter produces a controlled deformation imprint and the resulting residual profile is fitted via inverse finite element modelling to recover the true stress–strain response and tensile-equivalent properties (Figure 1(e)). A final validation campaign examined the Pareto-selected schedules using conventional tensile testing, hardness measurements, scanning electron microscopy (SEM), electron backscatter diffraction (EBSD), and correlation analysis against the model predictions.

2.2. Bayesian optimisation protocol

In this study, MOBO is used as a sequential decision-making framework that learns from prior experiments and proposes new heat treatment conditions expected to improve the trade-off between strength and energy

consumption. Conceptually, MOBO builds probabilistic surrogate models for each objective, quantifies uncertainty, and uses this information to efficiently search for Pareto-improving schedules without exhaustively sampling the full design space.

The MOBO protocol was initialised using a seed dataset generated from a Latin hypercube sampling (LHS) strategy. This provided well-dispersed starting conditions across the four-dimensional design space defined by solution-annealing and artificial-aging temperatures and hold times. Five initial points were selected within the predefined bounds (Table 1).

After the seed, the MOBO protocol proceeds iteratively: (i) training Gaussian process regression surrogate models on the available data to capture relationships between the four process variables (T_{SA}, t_{SA}, T_{AA}, t_{AA}) and the two objectives (UTS_{PIP} and E_C), yielding predictions and uncertainties over unexplored conditions; (ii) applying an Upper Confidence Bound (UCB) acquisition to balance exploration and exploitation using a single-alternating objective strategy, where one objective is optimised at each iteration while alternating between UTS_{PIP} and E_C, enabling balanced improvement of both objectives under a limited number of sequential experiments; (iii) enforcing a distance-based diversity filter to avoid repeated or closely clustered selections; (iv) experimentally evaluating the selected candidate and appending measured UTS_{PIP} and estimated E_C to the training set; and (v) repeating from Step (ii), continually refining the surrogates and updating the acquisition until the growth of the Pareto front (PF) becomes negligible.

Further implementation details, data processing steps, and code are provided in the Methods section, Supporting Information (sections S1.1 and S1.2), and the associated GitHub repository.

3. Results

3.1. Pareto front evolution

The Pareto front refers to the set of non-dominated solutions for which no objective can be improved without degrading the others. In Figure 2(a), the evolution of the PF during the MOBO campaign illustrates the progressive expansion of this trade-off boundary between UTS_{PIP} and E_C. Hypervolume (HV), an indicator of PF performance, is computed in the (UTS_{PIP}, $1/E_C$) objective space using a rectangular decomposition with a fixed reference point r = [200, 0] and provides a single scalar that increases as the PF moves into higher-strength and lower-energy regions (Figure 2(b)). The five seed samples (Table 1, samples S01–S05) provided broad coverage of the design space and defined the initial envelope. For example, Sample S01 required 5.13 kWh for 1125 MPa UTS_{PIP}, whereas Sample S05 consumed only 1.83 kWh and reached 1387 MPa. Subsequent fifteen MOBO iterations (Table S4, samples S01–E15) expanded the front as new candidates were explored. By iteration PF04, UTS_{PIP} and E_C reached 1451 MPa at 2.40 kWh, increasing the HV from the seed-achieved 649.64–676.46.

Later iterations refined both extremes: sample PF07 achieved 1045 MPa at the lowest energy input (1.40 kWh), while sample PF09 delivered the highest strength (1493 MPa) at 2.47 kWh, raising the hypervolume to 831.66. As shown in Figure 2(c), the steepest hypervolume growth occurred between iterations PF04–PF09, after which improvements diminished, indicating near-convergence. Beyond iteration 13, samples E13–E15 yielded consistent UTS_{PIP} values above 1438 MPa at intermediate energy levels (2.6–3.3 kWh), stabilising the hypervolume around 849. Figure 2(d) presents the heat-treatment profiles of the five Pareto-optimal samples, showing the corresponding solution annealing (T_{SA}, t_{SA}) and artificial aging (T_{AA}, t_{AA}) conditions that define the final Pareto set. The resulting Pareto front (Figure 2(e), Table 2) represents stable, non-dominated solutions with UTS_{PIP} values approaching 1500 MPa and energy consumption reduced by up to 60% compared to the most intensive seed schedules.

In addition to the observed HV growth, Figure 2(f) includes a retrospective hypervolume baseline to contextualise the efficiency of the acquisition-driven ordering. The baseline was constructed by fixing the first five seed experiments in their original order and recomputing cumulative HV in the objective space (UTS_{PIP}, $1/E_C$) using the same reference point $r = [200, 0]$. The 15 MOBO-acquired experiments were then randomly permuted 1000 times to generate baseline trajectories; at each post-seed iteration, the baseline mean and ±1σ envelope were computed and compared against the actual MOBO trajectory. By construction, all trajectories share the same initial HV (649.64) and converge to the same final HV (849.13), since the evaluated experiment pool is identical. Therefore, this comparison isolates ordering efficiency (how quickly HV is accumulated) rather than endpoint performance. Relative to random

Table 1. Seed dataset generated using Latin hypercube sampling with estimated energy consumption (E_C) and ultimate tensile strength (UTS_{PIP}) estimated by PIP.

ID	T_{SA} (°C)	t_{SA} (min)	T_{AA} (°C)	t_{AA} (min)	E_C (kWh)	UTS_{PIP} (MPa)
S01	1162	99.2	814	303.45	5.1309	1125.3
S02	1035	94.8	635	442.16	3.4689	1120.7
S03	1101	72.4	677	198.57	3.1452	1221
S04	1081	40.9	790	357.36	3.4559	1316
S05	964	27.3	707	132.44	1.8267	1386.7

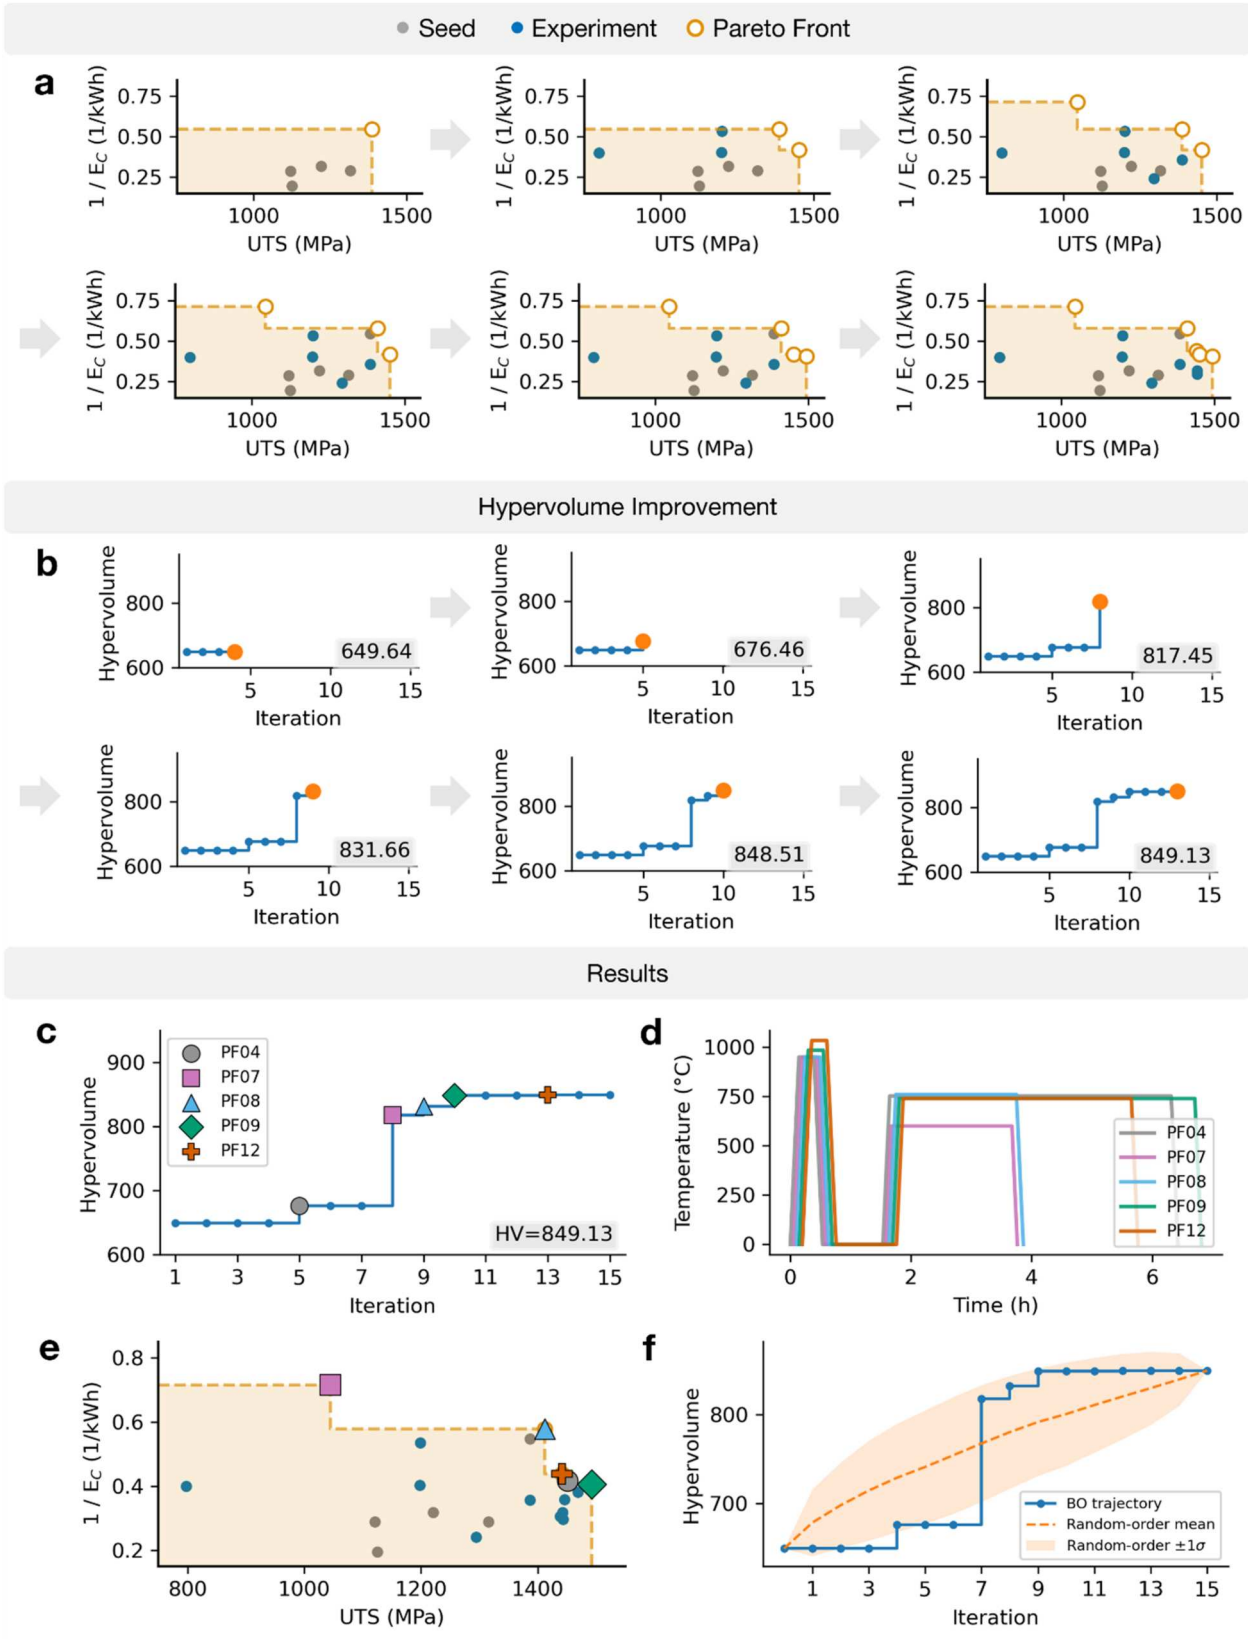

Figure 2. Pareto front evolution and results. (a) Evolution of the Pareto front over six representative iterations, showing the progressive expansion of the trade-off between ultimate tensile strength and energy consumption. (b) Hypervolume growth with iteration number, indicating Pareto improvement and convergence toward saturation. (c) Full hypervolume with marked Pareto-optimal samples. (d) Pareto-optimal heat-treatment schedules. (e) Distribution of all tested conditions in the two-objective space, highlighting initial points, BO-suggested candidates, and final Pareto-optimal solutions. (f) A retrospective baseline comparison of cumulative hypervolume in (UTS_{PIP}, $1/E_C$) using reference point $r = [200, 0]$.

Table 2. Summary of Pareto-optimal heat-treatment schedules identified by the MOBO protocol, showing SA and AA parameters (temperature and time), E_C, and UTS_{PIP}.

ID	T_{SA} (°C)	t_{SA} (min)	T_{AA} (°C)	t_{AA} (min)	E_C (kWh)	UTS_{PIP} (MPa)
PF04	950	15	753	280	2.4	1451.0
PF07	950	15	600	120	1.4	1044.7
PF08	950	15	760	120	1.7	1411.7
PF09	985	15	740	294	2.5	1492.7
PF12	1034	15	740	227	2.3	1441.3

ordering, MOBO was slower during the first six post-seed iterations; however, at iteration 7, it identified a high-value condition that produced a marked hypervolume jump (817.45) and moved the MOBO trajectory above the random order mean (767.53). MOBO then remained above the random baseline mean through the subsequent iterations before all trajectories converged to the same final hypervolume. This indicates that MOBO's benefit in this campaign is primarily a mid-to-late-stage ordering advantage. Thus, the main benefit of the acquisition strategy here is faster accumulation of useful Pareto improvements during the campaign, while the final non-dominated solution set remains unchanged because it is determined by the full set of evaluated experiments.

Table 3 summarises campaign-level performance metrics and the final surrogate accuracy under leave-one-out cross-validation (LOOCV). The final hypervolume reached 849.13, corresponding to a 1.3× expansion of the strength–energy trade-off relative to the initial seed value of 649.64. The energy-saving and UTS_{PIP}-gain factors of 2.18 and 1.21 indicate that the MOBO process identified schedules that reduced estimated energy consumption while improving UTS_{PIP} relative to the baseline seed designs. A hit ratio of 0.33 indicates that one-third of MOBO-acquired experiments contributed to the final Pareto front, consistent with exploratory sampling under a constrained experimental budget.

Table 3. Summary of MOBO campaign performance and final surrogate accuracy.

Campaign performance	
Initial seed points	5
MOBO experiments	15
Final hypervolume	849.13
Energy saving factor[a]	2.18
UTS_{PIP} gain factor[b]	1.21
Hit ratio[c]	0.33
Temporal cost[d](iterations)	8
Final surrogate accuracy (LOOCV)	
UTS_{PIP} RMSE (MPa)	45.25
UTS_{PIP} R^2	0.934
Energy RMSE (kWh)	0.48
Energy R^2	0.682

[a]Energy saving factor is computed as the best energy objective value observed over all evaluated points divided by the average energy objective value of the initial seed set.
[b]UTS_{PIP} gain factor is computed as the maximum UTS_{PIP} observed over all evaluated points divided by the average UTS_{PIP} of the initial seed set.
[c]Hit ratio is the fraction of MOBO-acquired experiments that appear on the final Pareto front.
[d]Temporal cost is the number of MOBO iterations after the initial seed design needed to first exceed 95% of the final hypervolume.

The final surrogate models exhibited strong predictive consistency under LOOCV, with $R^2 = 0.934$ for UTS_{PIP} and $R^2 = 0.682$ for E_C, and corresponding RMSE values of 45 MPa and 0.48 kWh. These results indicate that, once trained on the full dataset, the surrogate captures the dominant structure of the response surfaces. Convergence was achieved after eight MOBO iterations beyond the initial seed design, defined as the point at which 95% of the final hypervolume was first exceeded, consistent with the observed plateau in hypervolume improvement.

GP surrogate models, paired with a UCB acquisition function, define the core of the optimisation workflow. UCB was selected as a pragmatic choice because it is simple to implement and reasonably robust to noise, typically delivering mid-range performance among standard acquisition functions without being overly sensitive to noise levels; however, it is known to favour early exploitation once a good region is identified rather than maintaining an aggressive exploration–exploitation balance [31], which in this study is intentional given the small experimental budget and the priority of rapidly converging to a mechanically solid, low-energy solution.

The protocol was first executed using PIP-derived tensile-equivalent properties as the mechanical objective and was subsequently validated through an independent second MOBO loop. Using this verified setup, the optimiser was applied to a hardness-based objective in a total of 16 additional heat-treatment experiments, comprising the same five initial starting points (S01–S05) followed by 11 BO-guided experiments (Table S5). The campaign was terminated once the experimental budget had been reached. In the hardness-based MOBO campaign, the resulting Vickers hardness values spanned approximately 314–457 HV, while furnace energy consumption ranged from about 1.45 to 5.13 kWh, leading to the identification of seven Pareto-optimal schedules that captured the trade-off between mechanical performance and energy use. When assessed against the PIP-driven optimisation, the hardness-based results exhibited a consistent relative ranking of favourable heat-treatment regions and a similarly shaped Pareto front. This agreement supports the robustness of the optimisation strategy, indicating that the presented protocol captures processing-property relationships rather than artifacts of a particular mechanical proxy. Further details on the PIP-based campaign, including Pareto front evolution and hypervolume convergence, are provided in the Supporting Information,

section S1.3 (Figures S3, S4). Details on the hardness-based campaign, including model validation, sensitivity analysis, microstructural characterisation, and mechanical properties of selected Pareto samples, are provided in section S1.4 (Figures S8–S14, Table S6).

3.2. Microstructure of pareto front samples

Electron Backscatter Diffraction revealed that all Pareto-optimised heat treatments preserved the columnar grain morphology typical of LPBF-processed Inconel 718 (Figure 3(a) and (b)). Grains remained elongated along the build direction, reflecting the solidification sequence inherent to layer-wise melting. The persistence of this structure after heat treatment indicates that the applied temperatures (950–1034°C) and times (15 min) were below the full recrystallization threshold. Relative to the as-built condition (Figure 3(c)), the optimised samples did not show meaningful texture evolution: the texture remained weak in all PF variants and stayed broadly aligned with the build direction, with only minor differences between conditions. Overall, the PF heat treatments altered precipitation behaviour without meaningfully changing the solidification-derived texture.

Pole figures of the 001 orientation (Figure 3(d)) show weak ⟨001⟩ alignment across all conditions (peak pole density = 2.31–2.75), indicating minimal texture evolution. Differences in ⟨001⟩ alignment among PF07, PF08, and PF12 are small and within a narrow range, suggesting that the observed texture variations are section-specific rather than heat-treatment driven. Grain aspect ratios remain comparable to the as-built condition, confirming that stress-relief and aging treatments preserve the overall columnar grain morphology, with no evidence of significant recovery, recrystallization, or texture modification under the tested conditions.

Misorientation analysis (Figure 3(e)) showed that the distribution of low-angle (< 5°) and high-angle (15–59°) boundaries in the heat-treated samples remained broadly similar to that of the as-built state. No systematic shifts toward higher fractions of either boundary type were observed across the Pareto set. The grain-size and aspect-ratio histograms in Figure 3(f) and (g) are therefore shown as qualitative descriptors of the mapped regions. Each condition was represented by a single EBSD map acquired from the sample interior (region of interest $\approx 2 \times 2$ mm), so the dataset is sufficient to show retained columnar morphology and the absence of obvious recrystallization for all PF samples. The absence of Σ3 twin boundaries across all samples further supports that recrystallization did not initiate under any tested condition [1].

Nevertheless, texture similarity is expected, as most Pareto-optimised schedules employed mild solution annealing parameters (between 950–1034°C for 15 min). This short duration reflects the MOBO optimisation outcome, which prioritised energy efficiency, as furnace operation at higher temperatures is significantly more energy-intensive than extended holds at lower temperatures. Such short solution annealing treatments are also below the conditions generally required to substantially modify the solidification-derived texture or initiate recrystallization in LPBF IN718, which is consistent with both the present EBSD observations and prior literature [4,5].

Scanning electron microscopy with an In-Lens detector revealed distinct phase transformations across the Pareto-optimised heat treatments and the as-built sample (Figure 4(a)–(f)). The as-built condition exhibited a fine cellular dendritic substructure aligned with the build direction, characteristic of rapid solidification in laser powder bed fused Inconel 718. This structure is composed of parallel dendritic cells arranged within columnar γ-phase grains (Figure 4(a)). The substructure consists of packets of dendritic cells with small misorientation angles ($2° < \theta < 5°$), producing a high-density dislocation network. These packets are separated by low-angle grain boundaries ($5° < \theta < 15°$) within larger grains. The dendritic orientation varies across packets, and intercellular boundaries contain Laves phase and MC-type carbides formed through Nb, Ti, and Mo segregation, which are strongly influenced by laser power and scanning speed [2].

As discussed, the solution-annealing step dominates overall energy consumption because it operates at a higher temperature and requires a longer ramp to reach the soak. By contrast, the artificial-aging step contributes more directly to the development of tensile strength, as it controls the precipitation of strengthening phases such as γ′ and γ″. Consequently, the optimiser tends to select SA conditions at the lower end of the explored temperature range and at the minimum soak time (15 min), which reduces E_C without significantly degrading UTS_{PIP}.

The Pareto-optimal heat treatments all start from the cellular–dendritic LPBF microstructure, with the segregation of Nb- and Ti-rich interdendritic regions and associated Laves/MC-type carbides in the as-built state. The following interpretation of the Pareto-optimal samples is based on SEM-observed precipitate morphology and distribution, supported by established phase-evolution trends reported for both conventionally processed and additively manufactured IN718. The observed precipitate populations indicate systematic differences in precipitate density, morphology, and

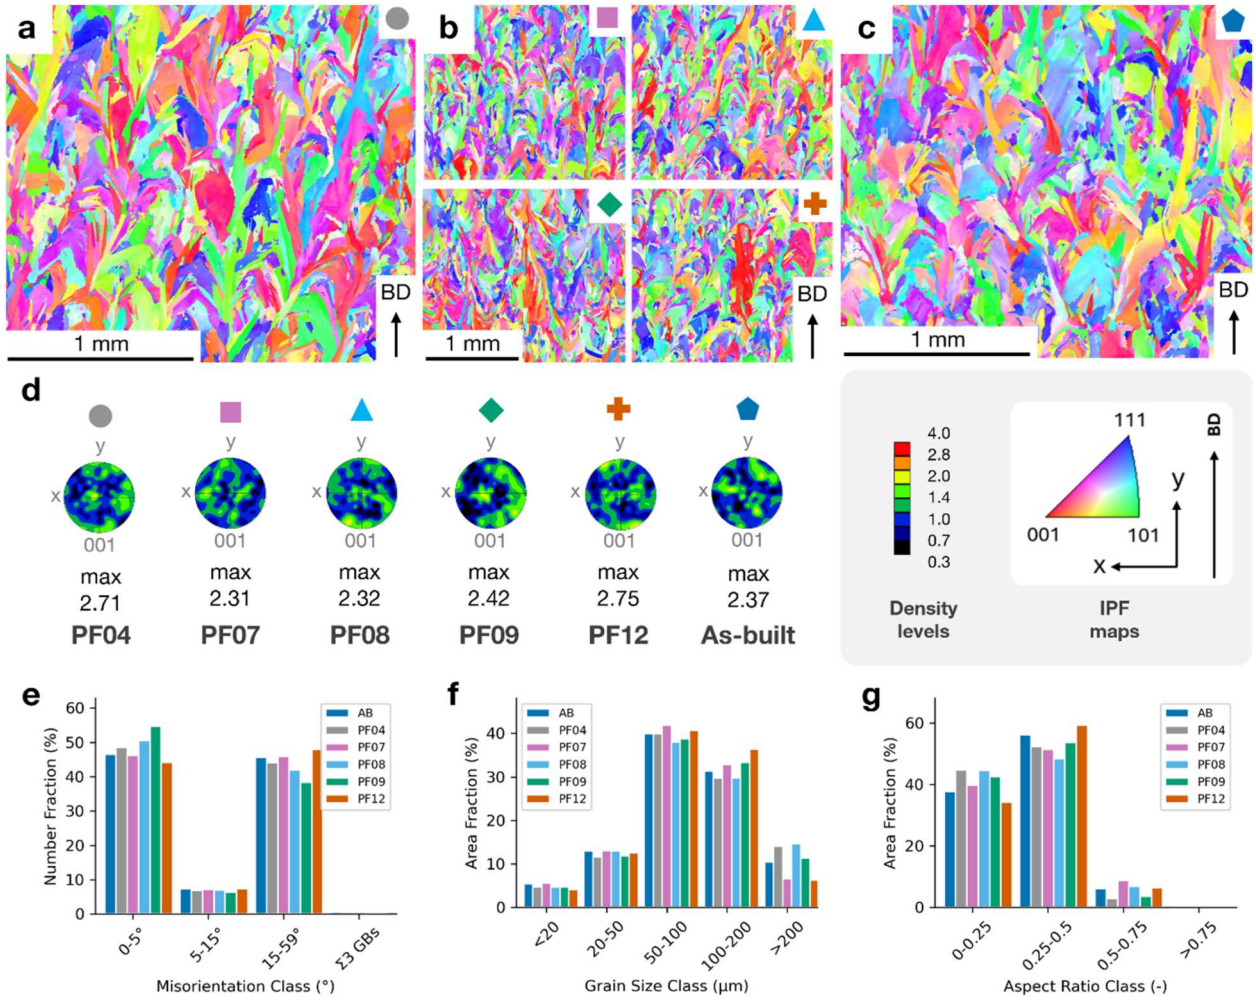

Figure 3. EBSD characterisation of Pareto front samples. (a) EBSD inverse pole figure (IPF) map along the build direction (BD), with the out-of-plane direction oriented vertically, showing elongated columnar grains in sample PF04. (b) IPF maps of additional Pareto front samples (PF07, PF08, PF09, PF12) demonstrating comparable grain morphology. (c) IPF map of the as-built sample. (d) {001} pole figures corresponding to panels (a)–(c), showing a dominant ⟨001⟩ texture aligned with BD; PF12 exhibits the strongest intensity (max = 2.75). Inset: IPF colour key. (e) Grain-boundary misorientation distributions highlighting the prevalence of low-angle (< 5°) and intermediate-angle (15–59°) boundaries. (f) Grain-size histograms for the mapped regions, showing that most detected grains fall within the 50–200 μm range. (g) Aspect-ratio histograms for the mapped regions, showing comparable grain elongation in the examined EBSD areas.

grain-boundary decoration across the Pareto-optimal conditions, consistent with differences in Nb redistribution during solution annealing and the subsequent balance between grain-boundary δ precipitation and fine intragranular γ′/γ′′ formation during aging. Solution annealing at 950°C for 15 min (PF04, PF07, PF08; Figure 4(b)–(d)) is below the δ-solvus and is expected to partially dissolve Laves while retaining a substantial fraction of Nb-rich regions that act as nucleation sites for the δ phase [32]. This is consistent with PF04 and PF08, which show dense, sub-micrometer needle-like precipitates distributed both within grains and along grain boundaries, together with finer globular features that likely correspond to large γ′/γ′′, small carbides, and remnants of the cellular substructure. Among these conditions, PF04 shows the densest and finest precipitate population, while PF08 exhibits the second-highest density together with more evident grain-boundary decoration. PF07, aged at 600°C, exhibits fewer fine precipitates but coarser needles and larger, high-contrast particles, which we interpret as a mixture of relatively coarse δ and stable carbides. In PF07, precipitation is also much less evident within grain interiors, suggesting a weaker intragranular strengthening contribution than in PF04 or PF08.

Increasing the SA temperature to 985°C (PF09; Figure 4(e)) promotes further dissolution and redistribution of Nb, leading to fewer intragranular precipitates and a stronger concentration of both needle-like and globular phases at grain boundaries after 294 min of aging at

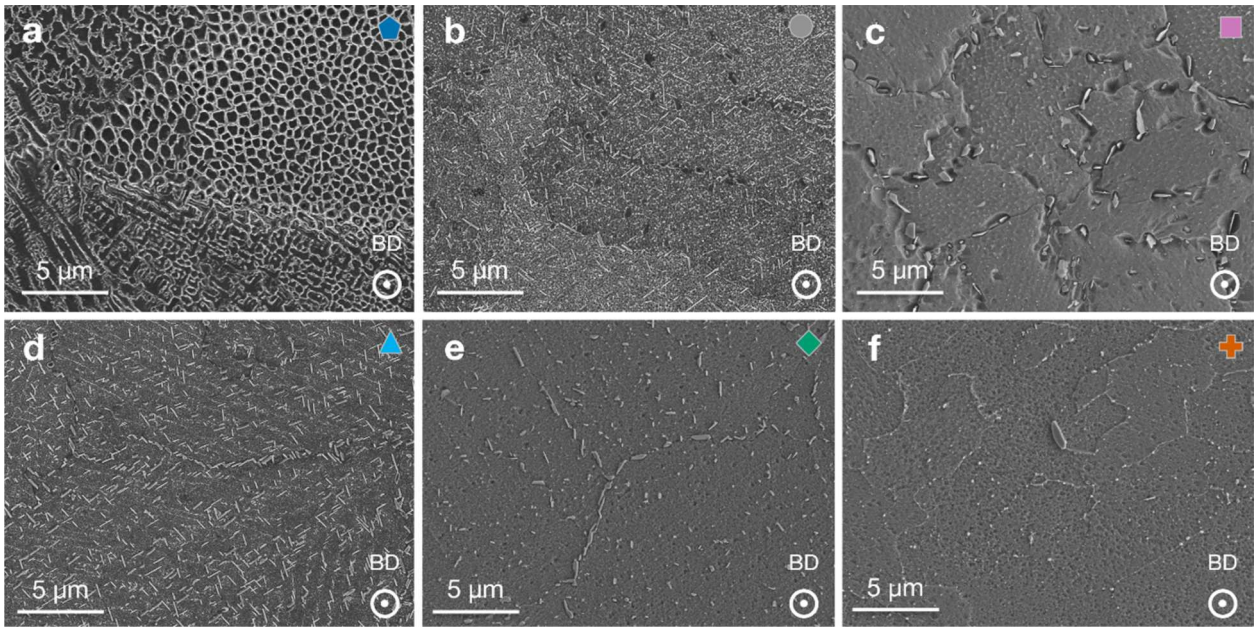

Figure 4. Microstructural characterisation of Pareto front samples. (a) As-built LPBF IN718. (b) PF04 (950°C/15 min + 753°C/280 min). (c) PF07 (950°C/15 min + 600°C/120 min). (d) PF08 (950°C/15 min + 760°C/120 min). (e) PF09 (985°C/15 min + 740°C/294 min). (f) PF12 (1034°C/15 min + 740°C/227 min). All SEM In-Lens micrographs were taken on sections perpendicular to the build direction and etched with Glyceregia.

740°C. Relative to PF08, PF09 shows a lower overall precipitate density, slightly coarser features, and much less precipitation within grain interiors, where fine precipitates remain mainly in regions associated with the former dendritic-cell structure. At the highest SA temperature (1034°C, PF12; Figure 4(f)), above the δ-solvus at ≈ 1010°C, the microstructure shows almost no needle-like precipitates, minimal evidence of residual cellular dendrites, and only sparse, sub-100 nm particles at grain boundaries. This is consistent with extensive dissolution of Laves/δ during SA and a dominant contribution from fine γ′/γ″ precipitates, which typically remain below SEM resolution in well-aged IN718, with characteristic sizes below 35 nm [33].

Taken together, these observations indicate a progressive reduction in the volume fraction and continuity of grain-boundary precipitates, particularly δ, as SA temperature increases. The δ phase is known to play a dual role in Inconel 718: a limited, discontinuous population along grain boundaries can pin boundaries, control grain size, and improve microstructural stability, thereby supporting strength and reducing notch sensitivity [34]. In contrast, fine intragranular γ′/γ″ precipitates are the primary strengthening phases in aged IN718, as they hinder dislocation motion and thereby increase yield and tensile strength. However, excessive δ along boundaries promotes microvoid nucleation during ductile fracture, degrades ductility, and when present as large particles, depletes Nb that would otherwise form strengthening γ″, which has adverse consequences for creep resistance [35, 36]. Several studies, therefore, argue for an 'optimal' δ fraction and morphology, where small, discontinuous boundary precipitates can impede crack growth at high temperatures without forming a continuous, embrittling network. Accordingly, the mechanical response reflects a balance between intragranular strengthening by γ′/γ″ and the competing effects of grain-boundary δ, whose increasing density and continuity promote boundary-controlled damage initiation.

Within this framework, PF04 and PF08, although mechanically strong, contain dense needle-like precipitates both inside grains and along grain boundaries, consistent with a relatively high fraction of δ plates. Such morphologies are more likely to approach the detrimental regime of boundary decoration discussed above. PF07 differs from PF04 and PF08 in that the precipitates are coarser and less numerous, with very limited precipitation inside grains, which is qualitatively consistent with a less favourable balance between intragranular strengthening and boundary precipitation. In contrast, PF09 and especially PF12 exhibit a much lower density of these features, with PF12 showing only sparse, sub-100 nm particles at grain boundaries and little evidence of residual cellular segregation. From a microstructural perspective, this trend suggests a transition away from dense needle-like grain-boundary precipitation toward conditions in

which boundary decoration is reduced and the strengthening contribution is expected to arise predominantly from finer intragranular precipitates below SEM resolution. PF09 and PF12 are therefore the most attractive options: they combine high UTS_{PIP} with a reduced fraction and continuity of coarse boundary precipitates that are most likely to embrittle the alloy under fatigue or creep loading.

Additional microstructural data, including optical microscopy (OM) and SEM micrographs of cross-sections parallel to the build direction, are provided in the Supporting Information, section S1.3 (Figures S6, S7).

3.3. Mechanical validation of pareto front samples

Mechanical validation of the Pareto-optimal samples was performed by comparing PIP-predicted performance with uniaxial tensile tests and hardness measurements to evaluate whether the PIP-based screening can capture the relative mechanical response of the selected conditions. Therefore, the PIP and tensile results are compared here primarily to determine whether the screening approach preserves the ranking and overall strength trends across the Pareto-optimal conditions. Vickers hardness was measured on the same coupons used for PIP, with the indented surface normal to the build direction. For tensile testing, standard round bars were heat-treated under the corresponding PF conditions, built vertically with their loading axis parallel to the build direction, and tested in the as-built surface condition after LPBF.

Figure 5 summarises and compares the results from both methods. UTS_{PIP}, yield strength (YS_{PIP}), and necking strain (NS_{PIP}) for the five Pareto-optimal samples are shown in Figure 5(a) for PIP. $UTS_{Tensile}$, $YS_{Tensile}$, and strain at break ($A_{Tensile}$) were obtained from the tensile curves and are shown in Figure 5(b) for uniaxial tensile tests. NS_{PIP} corresponds to the strain at the onset of tensile instability as defined by the Considère criterion on the reconstructed true stress–strain curve, i.e. the point where the work hardening rate equals the true stress, before significant localised thinning develops [37]. NS_{PIP} is therefore expected to be much lower than $A_{Tensile}$, because it marks the onset of necking rather than final fracture and should not be interpreted as a direct surrogate for tensile ductility. However, only UTS_{PIP} entered the MOBO objective together with E_C, whereas NS_{PIP} is included here only as a post hoc comparison metric to assess the limits of PIP in describing deformation behaviour. PF12 is a clear outlier: its tensile curve exhibits an extended, almost flat plateau after yielding, which is consistent with dynamic recovery and recrystallization maintaining an approximately constant flow stress during continued plastic strain [38]. In such cases, a large fraction of ductility is accumulated in the plateau regime, and this deformation is assumed to be poorly captured by NS_{PIP}, leading to a strong underestimation of ductility in PIP for PF12. By contrast, PF09 also exhibits an extended plateau but does not show a comparable discrepancy between NS_{PIP} and $A_{Tensile}$. This may reflect differences in the underlying deformation mechanisms, with PF09 governed by heterogeneous, boundary-controlled plasticity that PIP captures reasonably well, whereas PF12 is dominated by homogeneous recovery associated with fine γ'/γ'' precipitates.

Representative stress–strain curves in Figure 5(c) show that PIP overestimates the ultimate tensile strength relative to tensile testing while preserving the ranking of PF samples. To quantify the agreement between both methods, UTS, YS, and NS/A were min–max normalised to a 0–1 scale for each property, and the absolute normalised difference between the two datasets was computed (Figure 5(d)). UTS differences remained small (0–0.18), confirming a strong correlation between PIP and tensile UTS. YS deviations reached up to 0.34, whereas ductility showed the largest discrepancies. The largest normalised ductility difference (0.64) occurs for PF04, where PIP predicts one of the highest NS_{PIP} values, while tensile testing places PF04 at the lowest $A_{Tensile}$. In contrast, PF12 shows a smaller normalised difference (0.17) because both methods rank PF12 among the least ductile conditions, even though the relative change in ductility is large. This is made explicit in Figure 5(e). Tensile $UTS_{Tensile}$ deviates from PIP by only +3% for PF07 but by up to −14% for PF04. For ductility, the tensile fracture strain is consistently higher than NS_{PIP}, with PF12 showing the strongest relative increase, consistent with the flat post-yield plateau in its tensile curve.

The systematic tendency of PIP to overestimate UTS by $\approx$ 10% can be rationalised by specimen geometry, loading configuration, and surface condition. The tensile coupons were tested in the vertical, as-built condition with a rough surface, whereas PIP was performed on ground sections taken from coupons normal to the build direction. Although the EBSD pole figures indicate only weak residual ⟨001⟩ alignment and minimal texture evolution across the PF conditions, some orientation-dependent yielding is still expected. In the previous study on LPBF IN718, specimens tested perpendicular to the build direction showed higher tensile strength and Young's modulus than specimens tested parallel to the build direction in both the as-built and heat-treated states, with UTS differences of 57–87 MPa and Young's-modulus differences of 31–36 GPa [39].

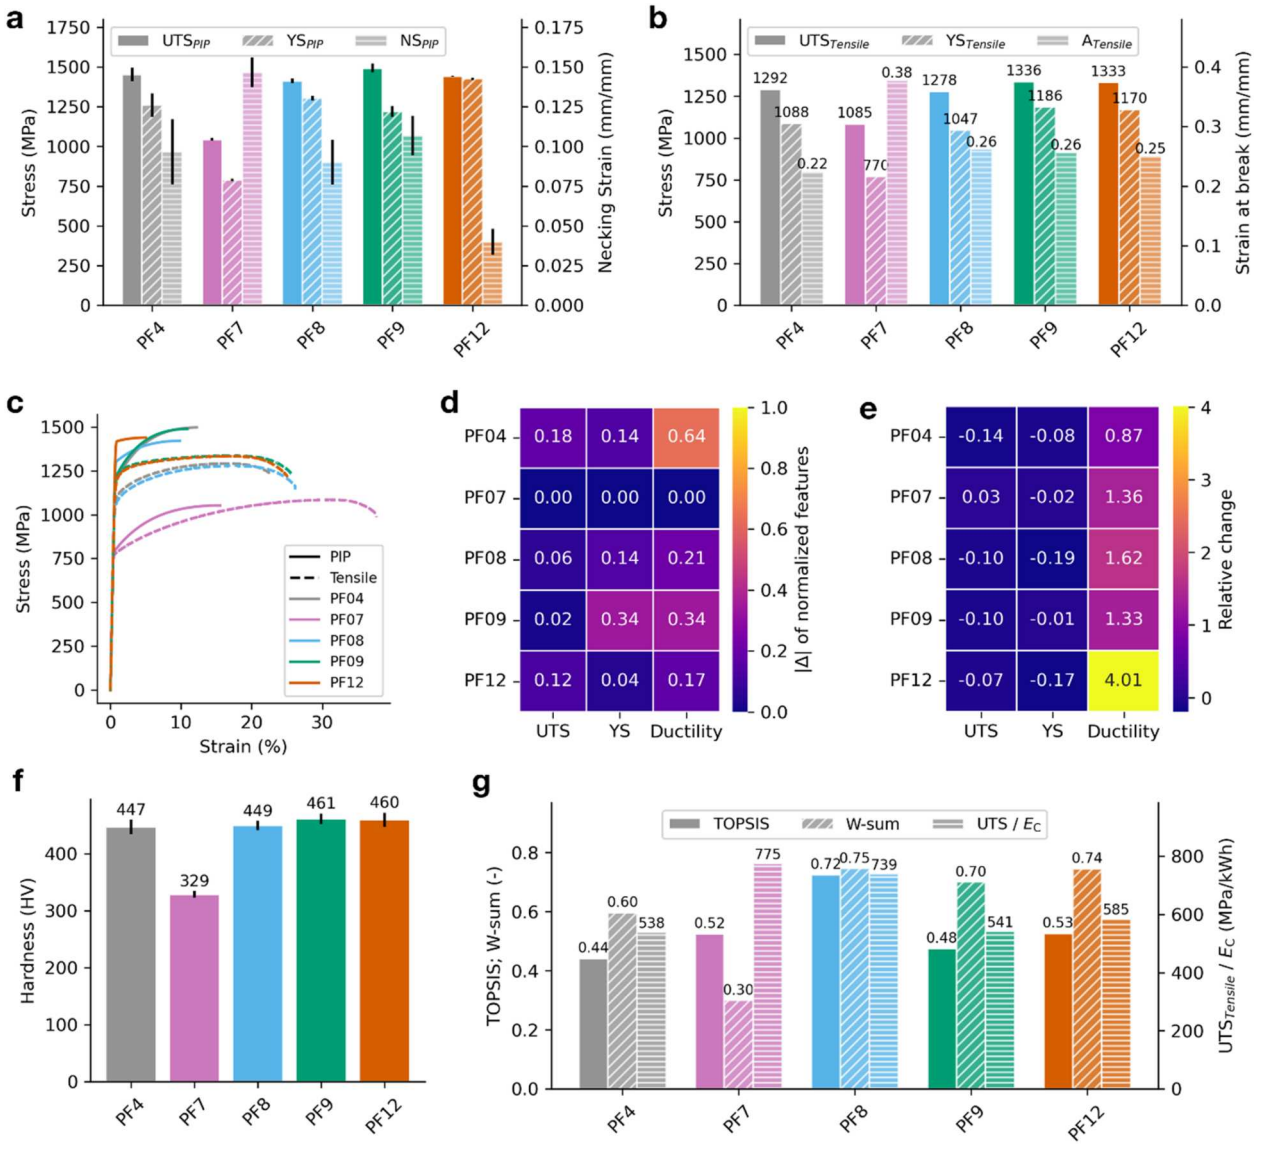

Figure 5. Mechanical validation of Pareto front heat treatments. (a) Mechanical properties of PIP-processed samples (mean ± SD). (b) Conventional tensile test results for the same conditions used to validate the PIP predictions. (c) Representative stress–strain curves from PIP and tensile tests. (d) Absolute difference, |Δ|, between PIP and tensile after separate min–max normalisation of each feature. (e) Relative change, (tensile − PIP)/PIP. (f) Hardness of the five samples (mean ± SD from ten indentations per sample). (g) Decision metrics for selecting a single compromise schedule on the Pareto front: TOPSIS closeness, weighted-sum score, and energy-normalised performance ($UTS_{Tensile}/E_C$) computed from tensile data for each condition.

In addition, the two methods differ substantially in sampled volume and surface condition: tensile testing integrates the response of the full gauge section, including surface-connected flaws and rare internal defects, whereas PIP probes a much smaller polished local region and is therefore less sensitive to defect-driven premature failure. In the present material, the spherical indenter radius (R = 1 mm) and characteristic grain size (typically 50–200 µm, Figure 3(f)) indicate that the PIP response reflects a locally averaged contribution from multiple grains rather than the orientation of a single grain. This interpretation is further supported by quantitative pore-area-fraction analysis from four representative cross-sections, which showed that the samples were highly dense (99.76% on average; Supporting Information Figure S1), with low pore fractions and no evidence of porosity differences large enough to explain the systematic PIP–tensile strength offset on their own. Because the main UTS ranking is preserved across the Pareto-optimal conditions, this offset does not affect the interpretation of the MOBO screening results. Combined with a few percent apparent strength reduction from the as-built rough surface (effective load-bearing area smaller than the nominal gauge area [40]),

a net difference of about ≈ 10% between PIP measured on sections normal to the build direction and tensile tests conducted parallel to the build direction is plausible, consistent with the deviations observed in Figure 5(e).

Although PIP ranks PF04 as highly ductile, tensile testing places it at the lowest fracture strain. Solution annealing at 950°C retains Nb-rich interdendritic regions and promotes dense, needle-like grain-boundary precipitates (Figure 4(b)). High Nb segregation and the resulting δ-like plates are known to stabilise grain boundaries but shift plastic deformation into grain interiors, promoting early strain localisation [41], and therefore reducing the uniform elongation measured in macroscopic tests. PIP, which samples a much smaller and more homogeneous volume, is less sensitive to this localisation, explaining the inflated NS_{PIP} estimate relative to the tensile $A_{Tensile}$.

In contrast, PF12, treated at 1034°C and aged, shows almost complete dissolution of the cellular segregation structure and of coarse Laves/δ, leaving only sparse, sub 100 nm boundary particles and a fine dispersion of γ'/γ'' that remains below SEM resolution (Figure 4(f)). This more homogeneous microstructure reduces the propensity for early boundary-controlled localisation. Because NS_{PIP} is derived from a localised deformation field around the PIP imprint and does not capture this long-range strain accumulation, it underestimates the true fracture strain for PF12. The different responses of PF04 and PF12 illustrate how NS_{PIP} is highly sensitive to the local balance between grain-boundary precipitates and intragranular strengthening phases. These PF04/PF12 mismatches therefore highlight the limitation of NS_{PIP} as a descriptor of bulk ductility, but they do not affect the suitability of PIP for screening candidates based on UTS in the present optimisation framework.

Vickers hardness measurements (Figure 5(f)) provide an additional check on mechanical performance across the Pareto set. Hardness followed the expected trends, with PF07 showing the lowest value (329 HV) and PF09 the highest (461 HV). Hardness levels in the 450–460 HV range are typical of fully aged LPBF IN718 [4], confirming that the PF aging schedules achieved near-complete precipitation strengthening.

Composite decision metrics in Figure 5(g) further illustrate the trade-offs along the Pareto front. For a $UTS_{Tensile}$-biased selection ($\omega\text{-}UTS_{Tensile} = 0.7$, $\omega\text{-}E_C = 0.3$), both TOPSIS and weighted-sum scalarization applied to tensile data identify PF08 as the preferred compromise, with PF12 as the second-ranked option. In contrast, a simple energy-normalised metric ($\omega\text{-}UTS_{Tensile}/\ \omega\text{-}E_C$) favours PF07 as the most energy efficient solution, with PF08 again ranking highly. The weighting $\omega\text{-}UTS_{Tensile} = 0.7$ was chosen to reflect typical high-value AM priorities, where satisfying mechanical qualification thresholds is the primary constraint and energy reduction is a secondary economic objective. The TOPSIS ranking remained unchanged across the range $\omega\text{-}UTS_{Tensile} = 0.6–0.8$, indicating low sensitivity to the specific weighting choice.

Microstructural considerations discussed earlier further differentiate the Pareto-optimal schedules. PF04 and PF08, while mechanically strong, retain dense needle-like precipitates within grains and along grain boundaries, consistent with elevated δ-phase fractions that can degrade ductility and fatigue resistance when present as continuous networks and deplete Nb from the matrix, reducing γ'' strengthening. PF07 exhibits coarser needles and carbide particles, further increasing susceptibility to crack initiation. In contrast, PF09 and especially PF12 show a much lower density and continuity of these phases, with PF12 displaying only sparse, fine boundary particles and minimal residual segregation. Accordingly, the Pareto front does not yield a single universally optimal condition but a set of viable schedules reflecting different priorities: PF07 favours energy efficiency, PF08 offers a strength–energy compromise, and PF09/PF12 represent mechanically robust and microstructurally safer choices where resistance to grain-boundary embrittlement and fatigue is critical.

4. Discussion

4.1. Surrogate models analysis

The performance of the GP surrogate models was evaluated using spatial uncertainty analysis and leave-one-out cross-validation (Figure 6). Uncertainty maps (Figure 6(a), (b), (d), and (e)) showed the lowest prediction variance near experimentally sampled regions, particularly around Pareto-optimal points ($\sigma = 0.153$ for SA and $\sigma = 0.202$ for AA). Higher uncertainty persisted at the boundaries of the design space where no experimental data were collected. This pattern reflects the adaptive acquisition strategy's focus on low-energy areas of the design space, which were repeatedly explored due to their higher expected utility.

The observed uncertainty pattern is consistent with how UCB controls sampling under a limited budget. In UCB, the exploration parameter β directly sets the trade-off between sampling high-uncertainty regions (large β) and exploiting regions with high predicted objective values (small β). Prior studies show that large, fixed β values can reduce exploitation bias but may allocate limited experimental resources to low-value exploratory points, while small β values accelerate

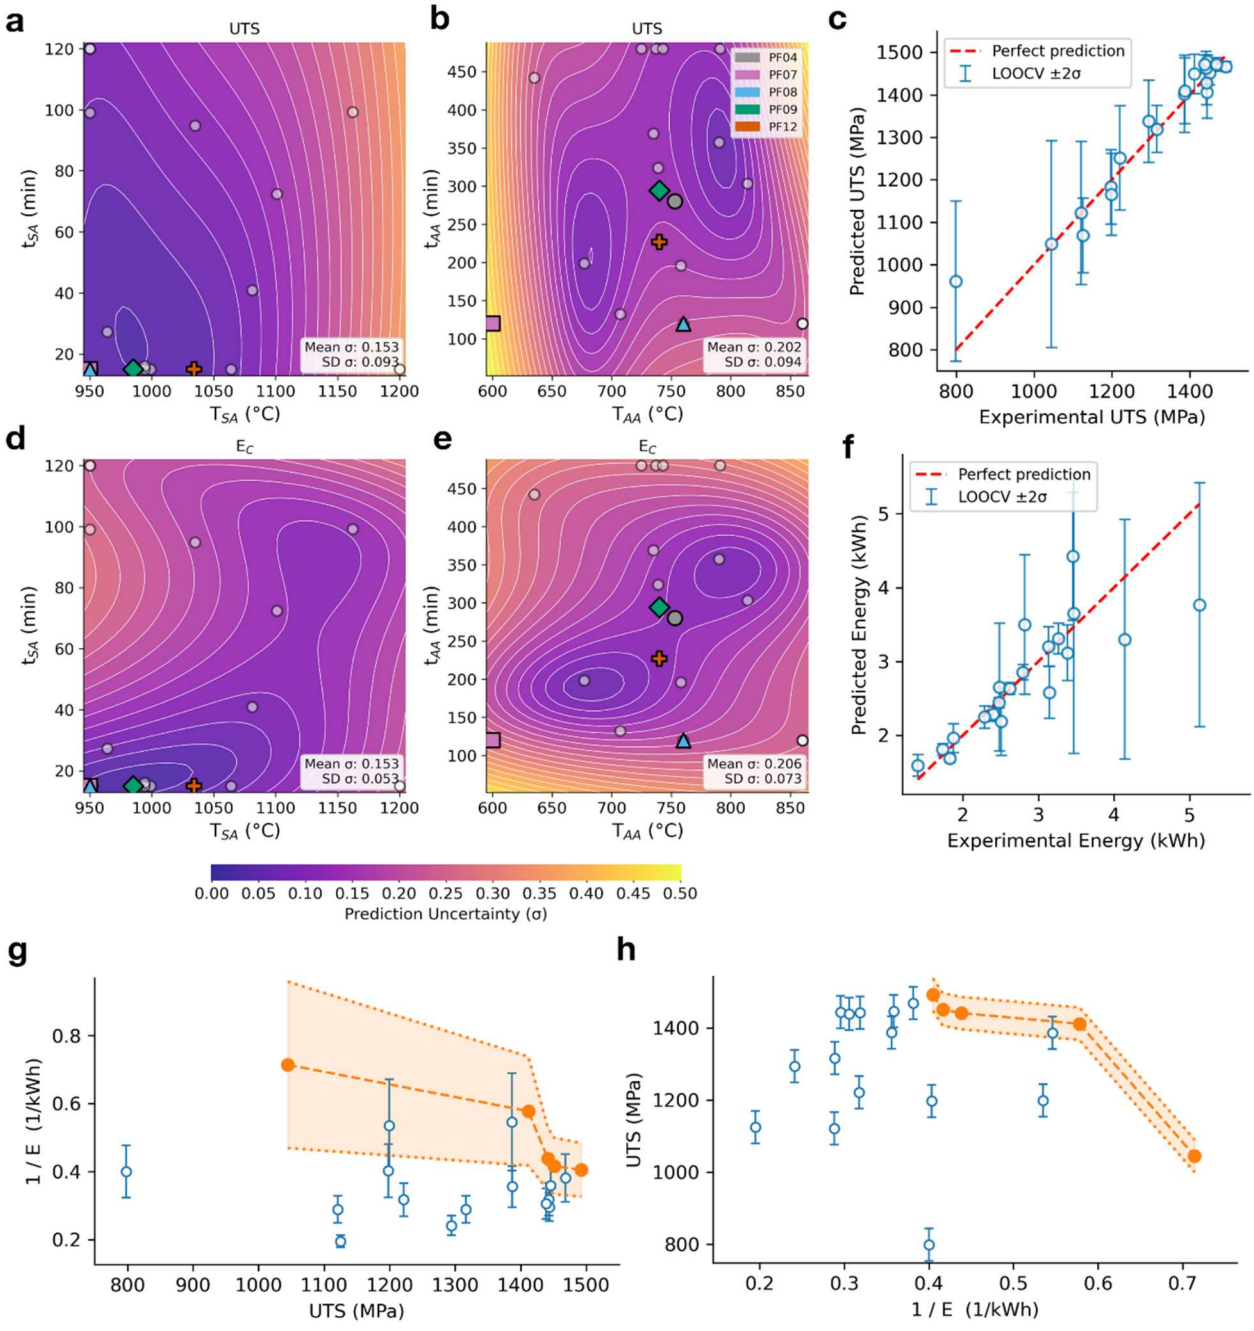

Figure 6. Model analysis and validation. Prediction uncertainty maps for ultimate tensile strength (a,b) and energy consumption (d,e) across the solution annealing and artificial aging treatment domains. Uncertainty is lowest near experimentally sampled regions (σ = 0.153–0.206) and increases toward untested boundaries. Leave-one-out cross-validation (LOOCV) of the surrogate models for UTS_{PIP} (c) and E_C (f). Predicted values closely match experimental data, with $R^2 = 0.934$ and RMSE = 45 MPa (3.5%) for UTS_{PIP}, and $R^2 = 0.682$ and RMSE = 0.48 kWh (17.1%) for E_C. Pareto-front uncertainty analysis using propagated surrogate errors: (g) $1/E_C$ with vertical error bars and confidence envelope, and (h) UTS_{PIP} with corresponding error bars, illustrate the stability of the Pareto-optimal solutions.

convergence yet increase the risk of premature stagnation in local optima in complex material spaces [42,43]. To avoid committing to a single, potentially suboptimal exploration level, randomised or adaptive β schedules have been proposed to maintain practical exploration early while reducing unnecessary exploration later, thereby improving robustness under tight experimental constraints [44,45]. In the present campaign, we intentionally operated in a 'good-solution-fast' regime: β was randomised within a moderate range, and the iteration count was capped at 15, which favours rapid identification of solid trade-offs but naturally leaves higher uncertainty in sparsely sampled, higher-energy regions of the design space. Candidate generation also

followed a single-alternating objective strategy, in which UCB optimisation targeted (UTS_{PIP}) and (E_C) in successive iterations, distributing sampling effort between objectives under the limited experimental budget and enabling progressive refinement of the trade-off without early bias toward one objective. A broader mapping of the full objective landscape would likely require either a more exploratory (larger/adaptive) β schedule or an explicitly cost/energy-aware acquisition term; both changes would increase the required number of iterations and were therefore not pursued here.

Additional safeguards were incorporated into the acquisition strategy to further reduce the risk of premature convergence. Acquisition optimisation employed multi-start initialisation (20 restarts, 2048 raw samples) to reduce sensitivity to local optima of the acquisition landscape. A distance-based diversity criterion was applied in the normalised parameter space by selecting candidates that maximise the minimum distance from previously evaluated conditions. When a proposed point fell within 0.15 of the normalised design space (≈ 15% of the parameter range), a bounded perturbation of ±(2–5)% of each parameter range was introduced to avoid clustering and promote continued exploration under the limited sequential experimental budget.

Cross-validation results (Figure 6(c) and (f)) quantified the predictive accuracy. For ultimate tensile strength (UTS_{PIP}), the GP model achieved an RMSE of 45 MPa (3.5%) and $R^2 = 0.934$, with most predictions within ±2σ. Deviations primarily appeared in the low-strength region, corresponding to under-sampled conditions. For energy consumption (E_C), the model showed larger residuals (RMSE = 0.48 kWh, 17.1%, $R^2 = 0.682$), concentrated above 3 kWh. These errors stem from the model's limited sampling of high-energy heat treatments, which were deprioritized by the acquisition function due to their low predicted performance-to-cost ratio.

To evaluate the influence of surrogate prediction error on the stability of the final Pareto front, uncertainty propagation was performed using the LOOCV RMSE values for both objectives (Figure 6(g) and (h)). The corresponding uncertainties were propagated to the Pareto space, producing error bars and confidence envelopes around the reported solutions. In Figure 6(g), the propagated E_C uncertainty (RMSE = 0.48 kWh) generates a sloped confidence band, with dominated candidates remaining below the upper bound of the front. In Figure 6(h), vertical uncertainty in UTS_{PIP} (RMSE = 45 MPa) widens the front locally, particularly in the intermediate trade-off region, while preserving the ordering of the Pareto solutions. These bounds reflect the non-uniform sampling inherent to sequential Bayesian optimisation, where sampling is concentrated in promising regions rather than uniformly across the design space. Despite higher uncertainty in sparsely sampled regions, no additional solutions are admitted to the Pareto front within the propagated bounds; thus, the Pareto set remains unchanged, and only the spacing between nearby solutions is affected.

Considered jointly, the Pareto uncertainty propagation (Figure 6(g) and (h)), the convergence of the hypervolume trajectory after eight iterations (Table 3), the random-order baseline comparison (Figure 2(f)), and the consistent results of the independent hardness-based replication campaign characterise the stability of the reported Pareto front within the constraints of the present dataset. The final Pareto set is defined by the full pool of evaluated experiments and is therefore independent of the specific acquisition sequence, with the replication campaign confirming a consistent relative ranking of favourable heat-treatment regions across a different mechanical proxy.

4.2. Parameter-objective relationships

The relationships between heat-treatment parameters and objective functions were evaluated using Spearman rank correlation analysis, complemented by a second-order linear model with two-way interactions and a non-parametric permutation importance analysis (Figure 7). This analysis is exploratory: the campaign was designed for sequential Bayesian optimisation rather than full-factorial inference, and with N = 20 observations and 10 model terms the available degrees of freedom are limited; *p*-values and effect sizes are therefore reported as descriptive indicators of relative effect magnitude.

Spearman's rank correlation analysis (Figure 7(a)) indicates that UTS_{PIP} exhibits weak to moderate monotonic associations with individual heat-treatment parameters ($\rho \approx -0.39$–0.42), with time-related variables showing stronger effects than temperature. Notably, UTS_{PIP} correlates negatively with solution time and positively with aging time. This behaviour is consistent with diffusion-controlled precipitation in IN718: strengthening is governed by the time-dependent evolution of coherent γ″ (and γ′) precipitates, for which precipitation occurs rapidly early in aging and then transitions to coarsening with diminishing strength gains, yielding sub-linear time responses [46]. The relatively weak temperature correlations are also expected within the constrained thermal window studied here. Peak γ″ precipitation rates occur around ~732–760°C, while dissolution occurs at substantially higher temperatures; similarly, γ′ forms over a comparable range and dissolves above ~843°C [47]. Within such 'active' bands, modest

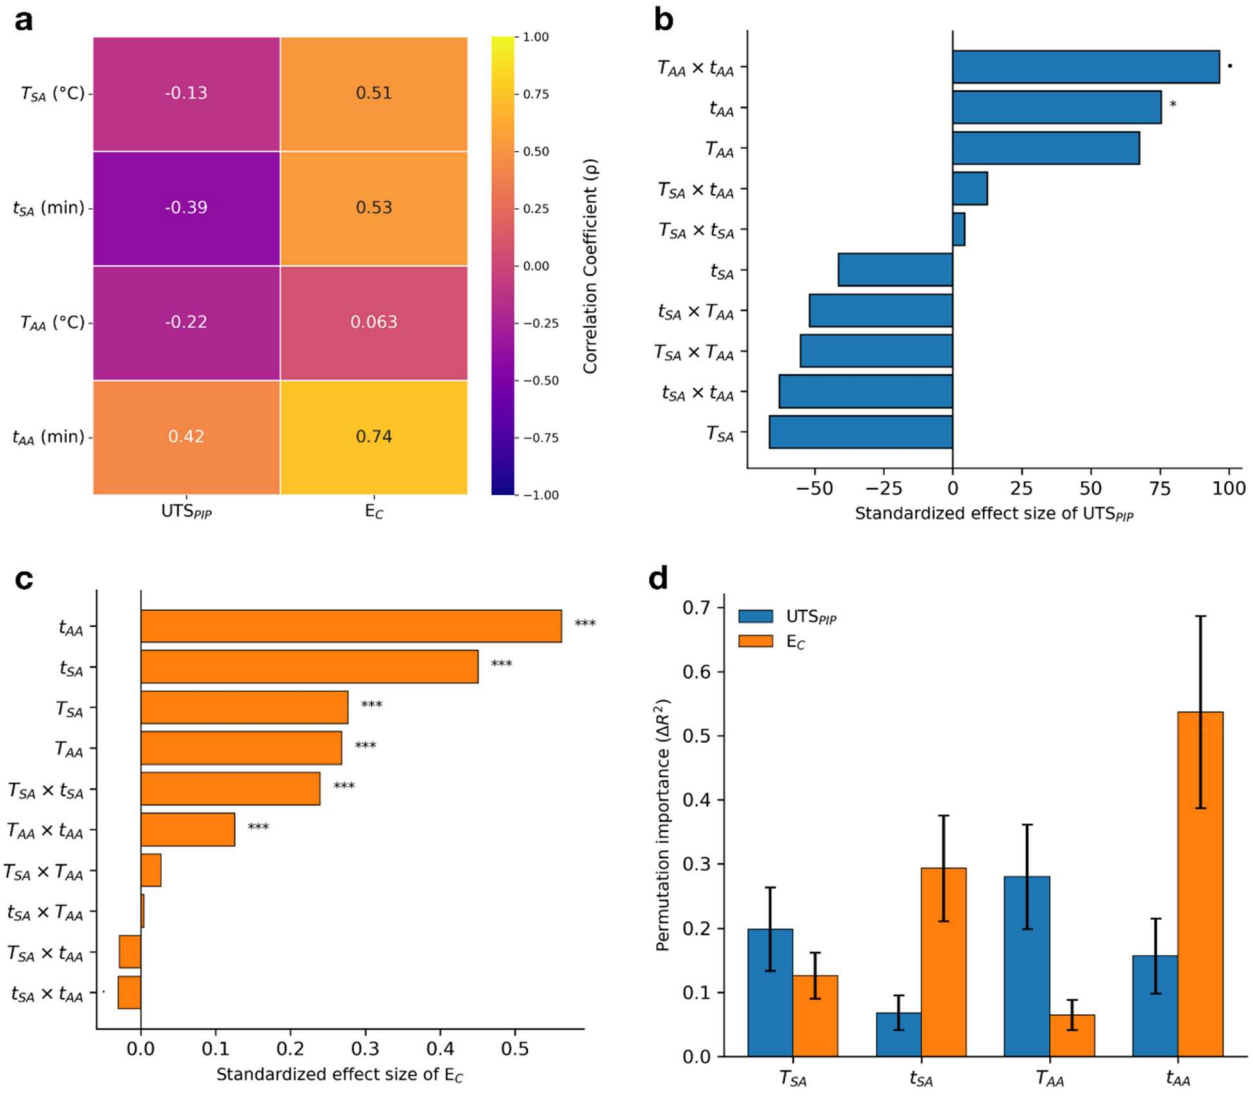

Figure 7. Parameter–objective sensitivity analysis. (a) Spearman's rank correlation heat map showing monotonic relationships between process parameters and objective functions. (b) Standardised β coefficients from a second-order linear model with two-way interactions for UTS_{PIP}; significance markers denote $p < 0.10$ (·), $p < 0.05$ (*), $p < 0.01$ (**), $p < 0.001$ (***). (c) Corresponding standardised β coefficients for E_C. (d) Permutation importance (mean decrease in R2± SD) from a Random Forest model (500 trees, 500 permutation repeats) for both objectives.

changes in T_{AA} mainly shift the time-to-peak rather than drastically change the final precipitate population. In addition, δ precipitation/dissolution introduces nonlinearity by competing with γ′′ through Nb consumption and by responding slowly to thermal exposure near the δ solvus, which can produce non-monotonic strength trends [48]. Complementary Pearson correlation and principal component analyses are provided in the Supporting Information, section S1.3 (Figure S5).

To quantify the individual and interaction contributions of heat-treatment parameters on both objectives, standardised β coefficients from the second-order linear model are shown in Figure 7(b) and (c). For UTS_{PIP} (Figure 7(b)), t_{AA} emerged as the only individually significant main effect ($\beta = 75.4$, $p = 0.04$), consistent with diffusion-controlled γ′′/γ′ precipitation being primarily governed by cumulative thermal exposure during aging. The $T_{AA} \times t_{AA}$ interaction term was marginally significant ($\beta = 96.4$, $p = 0.09$), indicating that the effect of aging time on strength depends on aging temperature, as expected from the coupled nucleation and coarsening kinetics of γ′′ within its active precipitation window. The overall weakness of remaining pairwise correlations also reflects the coupling among parameters in two-stage schedules (e.g. T_{SA} with t_{SA} and T_{AA} with t_{AA}), limiting the interpretability of single-variable metrics.

For E_C (Figure 7(c)), all four main effects were highly significant ($p < 0.001$), as were the $T_{SA} \times t_{SA}$ and $T_{AA} \times t_{AA}$ interaction terms, reflecting the near-deterministic

relationship between thermal cycle parameters and furnace duty-cycle energy demand ($R^2 = 0.998$). In contrast to UTS_{PIP}, energy consumption responded more strongly to all four parameters, as expected from furnace duty-cycle physics; the comparatively weaker contribution of T_{AA} to E_C relative to the hold times reflects both the narrow aging temperature window and the lower incremental power demand at aging temperatures compared with solution annealing.

As a nonparametric check that does not rely on linearity assumptions, permutation importance from a Random Forest model is shown in Figure 7(d). For UTS_{PIP}, T_{AA} and T_{SA} were the two most important features, indicating nonlinear effects related to δ and γ″ evolution that are not fully captured by the linear β coefficients. For E_C, t_{AA} and t_{SA} were the most important features, consistent with hold time being the main contributor to furnace energy consumption.

Across all three analyses, the aging parameters (T_{AA}, t_{AA}) are the main drivers of UTS_{PIP}, whereas the hold times (t_{AA}, t_{SA}) dominate E_C. The linear and nonparametric analyses are broadly consistent, with permutation importance additionally capturing nonlinear effects not reflected by the β coefficients. For E_C, all methods show the same trend: longer hold times increase energy consumption. Overall, Figure 7(b)–(d) points to the same parameter ranking and supports the interpretation used here.

4.3. Heat treatment benchmarking

The response of LPBF Inconel 718 to heat treatment is governed by the time-dependent precipitation of strengthening γ′ ($Ni_3(Al,Ti)$) and γ″ (Ni_3Nb) phases, which collectively control the overall mechanical properties. While precipitation occurs within a relatively narrow temperature window (approximately 600–900°C), the evolution of precipitate size, number density, and spatial distribution is highly sensitive to dwell time, resulting in mechanical properties that depend more strongly on cumulative thermal exposure than on temperature alone [49–51]. Above 650°C, γ″ becomes metastable and progressively transforms into δ, depleting Nb from the γ matrix and reducing γ″ strengthening. Extended hold times thus promote δ stabilisation at the expense of metastable phases, a process amplified by LPBF-specific features such as Nb microsegregation, high dislocation density, and retained Laves and MC-type carbides. CALPHAD analyses and TTT data indicate substantial overlap between γ″ and δ precipitation domains, so prolonged dwell times can induce phase evolution similar to that caused by higher temperatures [52]. Consequently, staged heat treatments exhibit strong multicollinearity, where the final microstructure reflects the integrated thermal history rather than individual variables.

To benchmark the optimised heat-treatment schedules, representative literature treatments were evaluated using the same performance metrics: $UTS_{Tensile}$ and E_C (Figure 8(a)–(c)). All E_C comparisons were recalculated using the same internal furnace model and are therefore intended for relative benchmarking under a common set of thermal-cycle assumptions. They should not be interpreted as absolute cross-study energy rankings, since actual power consumption depends on furnace-specific characteristics such as heating strategy, insulation, atmosphere, and load condition. The comparison of UTS and E_C shows that the Pareto-optimal schedules define the leading edge of the trade-off envelope, achieving comparable or higher strength at substantially lower E_C than the treatments reported in the literature. While some literature schedules reach high $UTS_{Tensile}$, they do so at significantly higher energy input: the highest reported $UTS_{Tensile}$ reaches 1471 MPa at 4.97 kWh, the most energy-intensive condition consumes 7.81 kWh for 1364 MPa, and another delivers approximately 1335 MPa while exceeding 5.4 kWh. By contrast, PF09 attains 1336.3 MPa at an estimated 2.47 kWh through a low-temperature, short-duration solution treatment followed by moderate aging, reducing energy consumption by roughly half for comparable strength. For reference, the standard aerospace heat treatment for IN718 (AMS 5663) corresponds to schedules S980-DA and S980-DA2 in the present benchmarking dataset: 980°C/1 h solution anneal followed by double aging (718°C/8 h + 621°C/8 h), yielding UTS of 1430–1471 MPa [50,53] and an estimated E_C of 4.97 kWh under the same thermal-cycle model used throughout. The relevance of this comparison is therefore not the absolute kWh value itself, but the fact that all schedules are evaluated within one consistent energy-cost framework. The Pareto-optimal schedules provide comparable UTS at ≈ 50% of this estimated energy input, although their lower solution temperatures and single-aging step imply different precipitate states and may require application-specific qualification. The corresponding strength–ductility comparison (Figure 8(b)) shows the Pareto-front solutions clustered near the optimal boundary, achieving competitive elongation despite reduced thermal exposure.

TOPSIS analysis (Figure 8(c)) places optimised schedules in the global decision space, with high-strength literature conditions favouring solutions that balance elevated $UTS_{Tensile}$ and moderate energy. Pareto front schedules cluster near the best-compromise region,

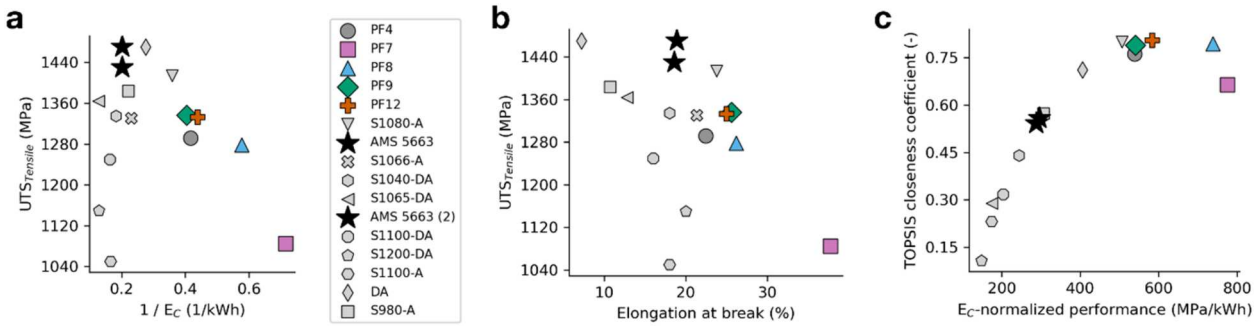

Figure 8. Benchmarking of optimised heat treatments against literature data. (a) Ultimate tensile strength ($UTS_{Tensile}$) plotted against reciprocal of estimated furnace energy consumption (E_C) for representative heat-treatment schedules reported in the literature. (b) $UTS_{Tensile}$ versus elongation at break for the same datasets. (c) Comparison of decision metrics used to select 'best-compromise': the TOPSIS closeness coefficient C_i ($UTS_{Tensile}$-biased weights $\omega_{UTS} = 0.7$, $\omega_E = 0.3$; larger is better) plotted against the energy-normalised performance $P_i = UTS_i/E_{C,i}$ (MPa/kWh; larger is better). Labels use S[ST]-A: S = solution treatment, ST = solution temperature (in °C), and A/DA = single/double aging. AMS 5663 marks AMS 5663 standard conditions; PFxx marks Pareto-optimal heat treatments for this study. All E_C values were recalculated using the same internal thermal-cycle model employed throughout this study; therefore, panels (a) and (c) should be interpreted as relative comparisons under common assumptions, rather than as absolute comparisons of measured energy consumption reported by different studies [1,39,50,53–57].

with PF12 showing the highest closeness coefficient (≈ 0.805) and PF07 achieving the highest energy-normalised performance (≈ 775 MPa/kWh). However, its lower absolute strength reduces its TOPSIS ranking. Although the optimised schedules are not universally optimal microstructurally, particularly for δ-phase morphology, these results demonstrate that, under the objective of maximising $UTS_{Tensile}$ while minimising energy input, the proposed schedules from the multi-objective Bayesian optimisation systematically outperform existing literature routes. The full benchmarking dataset, including heat-treatment schedules, is provided in the Supporting Information, Section S1.5.

The results indicate that the effective optimisation of heat-treatment processes for additively manufactured alloys depends on coordinated choices of (i) design-space bounds, (ii) objective definitions, and (iii) experimental cadence. Bayesian optimisation samples only within the admissible domain; thus, bounds should be selected to include the physical mechanisms that plausibly control performance (e.g. precipitation, dissolution, or phase instability), while excluding infeasible or application-irrelevant regimes that would waste the experimental budget. For established alloys, literature-informed limits and phase-stability constraints provide a rational starting point. Here, the design space and workflow were intentionally scoped so that each heat-treatment condition could be executed, mechanically screened, and reintegrated into the MOBO loop within one working day, reflecting realistic throughput under a constrained experimental budget.

Beyond this specific case, the workflow is compatible with a tiered (multi-fidelity) experimental strategy: rapid screening measurements enable dense exploration, while higher-fidelity tests are reserved for Pareto-selected conditions. In this study, PIP functioned as a strength-focused screening metric, allowing ranking and microstructural inspection on the same small coupon without dedicated tensile specimens during optimisation. Conceptually, the MOBO formulation is agnostic to the specific metric and can be extended to alternative objectives (e.g. ductility, fatigue indicators), wider heat-treatment windows, and other additively manufactured material systems. The same framing can also be expanded from post-processing to include process parameters (e.g. laser power, scan speed, hatch spacing), enabling coupled optimisation of processing and heat treatment within a unified decision space. More broadly, the present study argues for treating AM heat treatment as a process-chain design problem rather than as a fixed postscript to fabrication. Metal AM still relies heavily on conservative post-processing practice and experimentally expensive qualification routes, while post-processing itself contributes substantially to manufacturing cost and lead time. Within that context, the value of the present framework is to make resource intensity an explicit optimisation target while retaining metallurgical interpretability and practical experimental cadence toward faster and more resource-aware post-processing development for AM alloys.

4.4. Limitations and transferability

Despite this generality, several factors in the present campaign constrain the direct transfer of the identified

Pareto schedules and should be considered when applying the workflow to other machines, parameter sets, or alloys. The present study is limited to IN718 produced by a single LPBF platform with one parameter set. While the MOBO workflow is transferable in principle to complex additively manufactured components, including architected geometries, the specific Pareto schedules identified here may not transfer directly to such systems, where section size, local heat flow, and geometry-dependent deformation, crack-propagation, and damage-tolerance mechanisms can alter the effective response and therefore require dedicated revalidation [54,55]. Energy consumption was estimated using a semi-empirical thermal-cycle model rather than direct metering; the model captures relative differences but may carry systematic bias in absolute kWh values. The alternating single-objective UCB strategy is a pragmatic heuristic rather than a joint multi-objective acquisition function (e.g. qEHVI). A lightweight alternating heuristic was chosen to prioritise implementation simplicity and robustness under a tight sequential experimental budget of $N = 15$; hypervolume- based multi-objective acquisitions were not benchmarked in this study.

PIP was employed as a strength-biased screening tool: its role in the workflow is to filter candidates for high UTS_{PIP}, which are then validated by tensile testing. Its practical value in this context is that it provides richer mechanical screening information than scalar indentation metrics while requiring only a very small cubic specimen, which is advantageous for high-throughput AM workflows involving many candidate compositions, process parameters, or post-processing conditions. PIP showed reliable $UTS_{Tensile}$ ranking but large deviations in ductility, particularly for PF12 and PF04. Accordingly, the present study does not claim that PIP can serve as a rapid screening tool for ductility, fatigue performance, or fatigue life without additional calibration and validation against those specific endpoints. Because the optimisation targeted UTS_{PIP} and E_C, but not ductility, PIP is used here only as a strength-focused screening proxy. Ductility-sensitive applications would therefore still require tensile validation. All mechanical testing was performed at room temperature, and no elevated-temperature tensile, fatigue, or creep data were collected; application-specific qualification remains necessary.

Phase identification relies on SEM morphology and published phase-stability data. Quantitative confirmation of γ'/γ'' and δ phases, including their volume fractions, would require TEM, XRD, or calorimetric analysis, which were beyond the scope of the present campaign.

To apply this workflow to a new alloy, furnace, or component class, (i) define heat- treatment parameter bounds using CALPHAD phase-stability predictions or literature solvus and phase-boundary data; for architected geometries, the admissible design space and validation route should additionally reflect geometry-specific thermal and mechanical constraints; (ii) select a rapid mechanical proxy aligned with the objectives (e.g. PIP for strength-focused screening, micro-hardness for rapid filtering, or miniature tensile tests when ductility is critical) and re-calibrate the proxy-to-tensile mapping for the new alloy and microstructural state; (iii) characterise the furnace energy behaviour for the target setup by measuring ramp-rate curves and duty-cycle fractions, or replace the present analytical estimate with direct power metering when available, since the E_C model is system-specific and must be recalibrated for transfer to another production environment; (iv) initialise the campaign with $\approx$ 5 Latin Hypercube samples spanning the bounded design space; and (v) perform 10–15 sequential BO iterations with an experimental cadence that enables one closed-loop iteration per working day, consistent with the 4-dimensional implementation demonstrated here.

5. Conclusion

This work establishes an energy-aware closed-loop workflow for heat-treatment optimisation of Laser Powder Bed Fused Inconel 718. Rather than treating heat treatment as a fixed post-processing step selected from inherited schedules, the study frames LPBF post-processing as a sequential decision problem in which mechanical performance and process resource intensity are optimised jointly through PIP-based strength screening, MOBO, and a calibrated furnace-energy model.

- Starting from five LHS seed points and adding 15 sequential MOBO experiments, the campaign identified a Pareto set of heat-treatment schedules spanning low-energy and high-strength regimes, with a 31% increase in hypervolume relative to the initial seed set. The optimised schedules achieved up to a 1.21× increase in UTSPIP and a 2.18× improvement in the energy objective (1/EC) relative to the seed-set averages.
- The optimised solutions converged toward short solution-annealing holds combined with aging conditions that balance precipitation strengthening against process energy. Microstructural observations are consistent with heat treatment acting mainly through redistribution of segregation products and evolution of precipitate populations, while preserving

the solidification-derived columnar grain morphology and weak residual texture under the tested conditions.

- Tensile validation confirmed that the Pareto-selected schedules provide meaningful strength–energy trade-offs but also showed that PIP should be interpreted here as a strength-focused screening tool rather than as a general surrogate for ductility-related behaviour. The PIP-based strength ranking was reliable for the optimisation loop, whereas ductility-sensitive features required post hoc tensile assessment.
- Gaussian process surrogates coupled with uncertainty-aware sequential search successfully guided the campaign toward non-dominated regions under limited data, showing that useful decision-making can be achieved without requiring perfect absolute prediction.
- More broadly, the study shows that energy-aware HT optimisation for LPBF IN718 can be executed as a data-efficient search problem. The combination of PIP as a rapid mechanical screening signal and MOBO as an uncertainty-aware decision framework delivers a Pareto set of schedules that are competitive in strength to conventional treatments while substantially reducing estimated processing energy.

The framework is transferable in principle to other alloys and production environments, provided that the mechanical screening proxy and furnace-energy model are recalibrated for the target system. Future work should extend the workflow toward explicit multi-fidelity optimisation, additional application-driven objectives (e.g. ductility, fatigue resistance, notch sensitivity), and plant-specific resource usage metrics.

6. Methods

Bayesian optimisation protocol

We parameterise each two-stage heat treatment by four decision variables

$$x = [T_{SA}, t_{SA}, T_{AA}, t_{AA}],$$

where T denotes temperature (°C) and t denotes hold time (min), bounded within the ranges specified in Section 2. The two objectives are to maximise ultimate tensile strength, UTS(x), and to minimise furnace energy consumption per heat-treatment cycle, E(x). For MOBO, the second objective is transformed to its reciprocal, such that the model operates on

$$f_1(x) = UTS(x), \; f_2(x) = \frac{1}{E(x)},$$

so that both objectives are treated as maximisation problems. All temperatures and times are used in their physical units, then are linearly rescaled to the unit hypercube $[0, 1]^4$ prior to modelling; outputs are standardised to zero mean and unit variance.

Initialisation of the search space was performed via Latin Hypercube Sampling (LHS), generating a well-distributed set of $n = 5$ starting points across the four heat-treatment parameters. Two surrogate models, one per objective, were developed using Gaussian process (GP) regression implemented in the BoTorch Python library [56]. Each surrogate is a *SingleTaskGP* model operating on double-precision tensors, with a Matérn-type stationary covariance kernel with automatic relevance determination (ARD) and a Gaussian likelihood to capture observation noise (Table S3). Model hyperparameters (length scales, output scale, and noise level) are learned by maximising the exact marginal log-likelihood using gradient-based optimisation provided by GPyTorch [57].

At each MOBO iteration, an Upper Confidence Bound (UCB) acquisition function is used to select one candidate heat-treatment schedule. The UCB is defined as

$$UCB(x) = \mu(x) + \beta\,\sigma(x)$$

where μ(x) and σ(x) are the GP posterior mean and standard deviation at x, and β controls the exploration–exploitation balance. In practice, β is drawn uniformly from the interval [1, 3] at every acquisition call to introduce stochastic exploration while maintaining a bias toward promising regions. The bi-objective nature is handled by alternating between the two surrogates: in a given iteration, one of the objectives is selected, its UCB is maximised, and the resulting candidate is evaluated experimentally, after which the roles are swapped.

Acquisition optimisation is performed in the normalised input space using BoTorch's *optimize_acqf* routine. We use a batch size of q = 1, *raw_samples* = 2048, quasi-random initial points, and *num_restarts* = 20 multi-start local optimizations within the unit hypercube [0, 1]4. The resulting candidate is then mapped back to the original parameter ranges. To reflect practical furnace control, the solution-anneal and aging temperatures are rounded to the nearest integer degree Celsius before execution. All other parameters remain continuous within their bounds.

To avoid re-testing near-duplicate conditions, a diversity filter is applied to the set of acquisition-generated candidates. For each candidate, the minimum Euclidean distance is computed in the normalised parameter space to all previously tested points. The final experiment is chosen as the candidate that maximises this minimum distance. If this distance falls below 0.15 (indicating

that the candidate lies too close to an existing experiment), a small random perturbation is added: each parameter is perturbed by 2–5% of its allowable range with a uniformly distributed sign, and the result is then clipped back to the admissible bounds. This procedure maintains novelty while ensuring that all recommended schedules are physically feasible and implementable in the furnace.

Non-dominated points are identified using BoTorch's Pareto mask on the transformed objectives. Hypervolume is calculated over the (UTS, 1/E) space using a rectangular decomposition with a fixed reference point r = [200, 0], and the resulting scalar is stored alongside each experiment as a measure of optimisation progress. Additional bookkeeping (e.g. a binary improvement flag, hypervolume history) is written back to the same database file after each iteration.

The optimisation was run for 15 iterations in total, corresponding to 15 sequentially selected experiments beyond the initial seed. No automatic early-stopping rule was imposed in the implementation; in this study, the campaign was terminated once the hypervolume increase per iteration became negligible and the experimental budget was exhausted.

All MOBO routines were implemented in Python using PyTorch, BoTorch, and GPyTorch for surrogate modelling and acquisition optimisation, NumPy and Polars for data manipulation, and Matplotlib/Seaborn for visualisation. Furnace power rating and parameter bounds were supplied through a JSON configuration file, ensuring reproducibility of the design space and energy model. Pareto-optimal points and hypervolume trajectories were exported as CSV files and figures to support further analysis and archival.

Fabrication and heat treatment

All specimens were fabricated using a SLM Solutions 280 2.0 laser powder bed fusion system (Nikon SLM Solutions GmbH, Germany) equipped with a 700 W IPG 1070 nm continuous-wave fibre laser. Gas-atomized Inconel 718 powder with a particle size distribution of 15–45 µm was used as the feedstock material. A high-purity argon environment (Argon 5.0) was maintained throughout the process to minimise oxidation. The builds were performed on an IN625 substrate ($100 \times 100 \times 25$ mm3) and incorporated processing conditions designed to mitigate thermal distortions, including low heat input, fine layer thickness, preheating of the build platform, and a stripe-scan strategy with a 67° rotation between consecutive layers. This scanning strategy promoted uniform heat distribution and mechanical homogeneity within the fabricated parts. The process parameters were selected according to manufacturer recommendations for dense IN718 and comprised a laser power of 300 W, scan speed of 1300 mm/s, focal beam diameter of 100 µm, layer thickness of 30 µm, hatch distance of 120 µm, build-platform preheat to 200°C, and a stripe-scan strategy (stripe length 7 mm, two contour borders) with a 67° hatch rotation between successive layers (Table S1). These settings correspond to a linear energy density of 0.23 J/mm and a volumetric energy density of 64 J/mm3, and have been shown to yield IN718 parts with a relative density >99.95% and high dimensional accuracy [1,58].

Following the LPBF fabrication process, the as-manufactured samples were heat-treated in two stages: solution annealing with air cooling, followed by artificial aging with furnace cooling. All the heat-treatment experiments were conducted using a FCF-4/160M (Czylok Sp. z o.o., Poland) single-phase chamber furnace equipped with a microprocessor-based PID temperature controller (MRT 4 regulator) and high-temperature Type B thermocouple.

Energy consumption

Energy consumption (E_C) for each two-stage heat-treatment schedule was computed using the furnace's rated power Prated = 2.6 kW and the empirical model of the ramp and soak durations:

$$E_{total} = P_{rated} \cdot (t_{ramp,1} + D_1 t_1 + t_{ramp,2} + D_2 t_2)$$

where $t_{ramp,i}$ and t_i denote the ramp and soak durations for stage i, and D_i is the empirically determined duty cycle fraction at that temperature (Figure S2). The reported E_C values represent the estimated per-run electrical energy computed from the furnace's rated power and empirical duty-cycle model for a single heat-treatment run (on a per-batch basis), including both stages and the full ramp/soak segments. This model provides a first-order estimate suitable for ranking schedules by relative thermal cost. It does not capture batch-loading effects, furnace aging, ambient temperature variation, or transient power spikes during ramp segments. The reported E_C values should therefore be interpreted as a thermal-budget proxy rather than metered energy consumption.

Microstructure characterisation

After completion of the heat-treatment cycle, the Pareto front representative samples were subsequently polished for characterisation. The preparation sequence followed standard metallographic protocols: specimens were mounted in conductive resin to expose both the XY and XZ planes and subjected to a grinding sequence using silicon carbide abrasive papers of progressively finer grit sizes (120, 240, 360, 800, and 1200), with thorough cleaning between each step. Fine polishing was then performed using a colloidal silica suspension to achieve a mirror-like surface finish. To reveal

microstructure, the samples were etched using Glyceregia (5 ml HNO3, 10 ml glycerol, 15 ml HCl), with etch times between 45–120 s, using a lint-free cellulose cloth to gently wipe the samples. Optical microscopy examination was conducted using an OLS4000 LEXT microscope (Olympus Corp., Japan). Porosity was quantified from stitched confocal cross-sectional images (Olympus LEXT OLS4000) using a Python workflow built with Pillow (PIL), NumPy, and SciPy (scipy.ndimage). Each image was binarized at a threshold of 160, the largest connected white region was selected as the sample area, pores were segmented as black pixels, and pore-area fraction was calculated as black pixels divided by total pixels within the sample. Following optical characterisation, scanning electron microscopy (SEM) was performed using a Zeiss ULTRA 55 field-emission SEM operating in In-Lens mode at 5 kV and working distance between 3 and 5 mm. Electron backscatter diffraction (EBSD) was conducted at 20 kV, 120 µm aperture, high-current mode, and the resulting data were analyzed using EDAX OIM Analysis 8 software.

Profilometry-based indentation plasometry

Profilometry-Based Indentation Plastometry (PIP) was used to determine tensile-equivalent stress–strain curves for heat-treated samples in accordance with ASTM E3499-25. Cylindrical specimens (∅ 16 mm × 100 mm) were ground on both sides using SiC papers with grit sizes from 60 to 2000 to ensure flatness and surface consistency. Each heat-treatment condition was evaluated using three PIP measurements. Tests were conducted with a 1 mm spherical indenter using the PLX-Benchtop system (Plastometrex, UK). The residual indent profiles were measured by the instrument's integrated optical profilometer and analyzed using the Corsica software package. The inverse finite element fitting procedure, implemented within the software, iteratively adjusted the plasticity parameters until agreement between the experimental and simulated indent profiles was achieved. The resulting parameters were converted into true stress–strain curves, from which ultimate tensile strength and necking strain were derived.

The necking strain (NS_{PIP}) is defined as the tensile-equivalent strain at the Considère point, identified on the reconstructed true stress–strain curve obtained by inverse FEM fitting of the residual indentation profile [37]. NS_{PIP} corresponds to the predicted onset of tensile instability and is distinct from the fracture strain measured in uniaxial tensile tests.

Mechanical testing

Vickers hardness measurements were carried out using a µHV Vickers hardness tester (Zwick Roell Group, Germany) at a load of 2 kg and a dwell time of 10 s. For each specimen, the hardness values represent the average of ten indentations arranged in a grid pattern, with each indentation manually measured using the Clemex software interface. The indentations were spaced 450–500 µm apart to prevent overlapping strain fields and account for microstructural heterogeneity within the sample.

Cylindrical tensile specimens (5 mm diameter) with a 20 mm-long parallel gauge section and as-built surface condition were manufactured via LPBF and then heat-treated according to Pareto front heat-treatment schedules. The specimens were subjected to uniaxial static tensile tests with a strain rate of 0.5 mm m^{-1} s^{-1} at room temperature in accordance with ASTM E8 using a servo-hydraulic machine (MTS 318.10, Instron Corp., USA) equipped with an extensometer (MTS 63253F-14, Instron Corp., USA). Instron Bluehill software was used for data acquisition and analysis. From these data, key mechanical properties, including Young's modulus, ultimate tensile strength, tensile strength at fracture, and strain at fracture, were determined.

To rank Pareto-front heat treatments in terms of their proximity to an ideal compromise between strength and energy, we computed a TOPSIS (Technique for Order of Preference by Similarity to Ideal Solution) closeness coefficient for each schedule [59]. The decision matrix comprised the energy consumption per cycle, E_C (cost criterion; to be minimised), and $UTS_{Tensile}$ (benefit criterion; to be maximised). Following the classical TOPSIS procedure, each column was vector-normalised,

$$r_{ij} = \frac{x_{ij}}{\sqrt{\sum_i x_{ij}^2}}$$

and then weighted as $v_{ij} = w_j\, r_{ij}$. The ideal best (A^+) and ideal worst (A^-) points were defined from the extrema of the weighted normalised matrix: for UTS, $v^+_{UTS} = \max_i v_{i,UTS}$ and $v^-_{UTS} = \min_i v_{i,UTS}$, while for E_C, $v^+_E = \min_i v_{i,E}$ and $v^-_E = \max_i v_{i,E}$. Euclidean distances to these points,

$$d_i^+ = \sqrt{\sum_j (v_{ij} - v_j^+)^2},\ d_i^- = \sqrt{\sum_j (v_{ij} - v_j^-)^2},$$

were used to compute the TOPSIS closeness coefficient:

$$C_i = \frac{d_i^-}{d_i^+ + d_i^-},$$

where a larger C_i indicates a solution closer to the ideal (high $UTS_{Tensile}$, low E_C). We used a $UTS_{Tensile}$-biased weighting, $\omega_{UTS} = 0.7$ and $\omega_E = 0.3$.

For comparison, a weighted-sum scalarization was also evaluated on the same two objectives. Both E_C and $UTS_{Tensile}$ were linearly scaled to [0, 1] using min–max normalisation across the Pareto set, yielding $\tilde{E}_{c,i}$ and $\tilde{UTS}_i$ for each schedule i. A single composite score

was then defined as:

$$S_i = \omega_{UTS}\tilde{UTS}_i + \omega_E(1 - \tilde{E}_{c,i}),$$

with $\omega_{UTS} = 0.7$ and $\omega_E = 0.3$. Higher Si values correspond to treatments that increase

$UTS_{Tensile}$ while reducing E_C in the normalised space.

Finally, an energy-normalised performance metric was calculated as a simple ratio of tensile strength to energy consumption:

$$P_i = \frac{UTS_i}{E_{c,i}},$$

with $UTS_{Tensile}$ in MPa and E_C in kWh, providing an 'MPa per kWh' measure of strength delivered per unit furnace energy.

Author contributions

CRediT: **Enjy Katary:** Conceptualization, Data curation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing; **Kinga Sekuła:** Data curation, Investigation, Writing – review & editing; **Dawid Krok:** Data curation, Investigation, Software; **Marcin Posmyk:** Data curation, Investigation, Visualization; **Michał Dusza:** Data curation, Investigation, Software; **Wojciech Matusik:** Conceptualization, Supervision, Writing – review & editing; **Konrad Gruber:** Conceptualization, Data curation, Funding acquisition, Investigation, Methodology, Project administration, Visualization, Writing – original draft, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research was supported by the National Science Centre (NCN), Poland, under the OPUS 27 program, project UMO-2024/53/B/ST8/00721. E.K. acknowledges support from the STRAINS Erasmus Mundus Master's programme.

AI-use declaration

During manuscript preparation, large language models (OpenAI's GPT-5.1, GPT-5.2) were used to improve prose clarity and style. All suggestions were reviewed, verified, and approved by the authors. No AI tools were used to generate scientific content, analysis, or conclusions.

Data availability statement

The data supporting the findings of this study will be made publicly available upon publication at https://10.5281/zenodo.18323683. The MOBO framework code developed in this study will be made publicly available upon publication at https://github.com/Enjy-Katary/am-ht-mobo.

ORCID

Enjy Katary http://orcid.org/0009-0004-9192-4352
Kinga Sekuła http://orcid.org/0009-0007-8771-0193
Dawid Krok http://orcid.org/0009-0002-4666-5182
Marcin Posmyk http://orcid.org/0009-0003-8356-6500
Michał Dusza http://orcid.org/0009-0000-3823-8944
Wojciech Matusik http://orcid.org/0000-0003-0212-5643
Konrad Gruber http://orcid.org/0000-0003-1357-2473

References

[1] Gruber K, Dziedzic R, Kuznicka B, et al. Impact of high temperature stress relieving on final properties of inconel 718 processed by laser powder bed fusion. Mater Sci Eng A. 2021;813:141111.
[2] Gallmeyer TG, Moorthy S, Kappes BB, et al. Knowledge of process-structure-property relationships to engineer better heat treatments for laser powder bed fusion additive manufactured inconel 718. Additive Manuf. 2020;31:100977.
[3] Gruber K, Stopyra W, Kobiela K, et al. Mechanical properties of inconel 718 additively manufactured by laser powder bed fusion after industrial high-temperature heat treatment. J Manuf Process. 2022;73:642–659.
[4] Hosseini E, Popovich V. A review of mechanical properties of additively manufactured inconel 718. Additive Manufacturing. 2019;30:100877.
[5] Volpato GM, Tetzlaff U, Fredel MC. A comprehensive literature review on laser powder bed fusion of inconel superalloys. Additive Manufacturing. 2022;55:102871.
[6] Fayed EM, Saadati M, Shahriari D, et al. Effect of homogenization and solution treatments time on the elevated-temperature mechanical behavior of inconel 718 fabricated by laser powder bed fusion. Sci Rep. 2021;11:1–17.
[7] Salgado-Lopez JM, Martinez-Franco E, Cruz-Gonzalez C, et al. Microstructure and microhardness evolution of additively manufactured cellular inconel 718 after heat treatment with different aging times. Metals. 2022;12(12):2141.
[8] Benzing JT, Derimow N, Kafka OL, et al. Enhanced strength of additively manufactured Inconel 718 by means of a simplified heat treatment strategy. J Mater Process Technol. 2023;322:118197.
[9] Dwivedi A, Khurana MK. Optimization of the heat treatment strategy for Inconel 718 superalloy produced by laser powder bed fusion additive manufacturing process. Mater Sci Technol. 2025;0:02670836251352011.
[10] Laleh M, Sadeghi E, Revilla RI, et al. Heat treatment for metal additive manufacturing. Prog Mater Sci. 2023;133:101051.
[11] Clyne TW, Campbell JE, Burley M, et al. Profilometry-based inverse finite element method indentation plastometry. Adv Eng Mater. 2021;23:2100437.
[12] Tang Y, Campbell J, Burley M, et al. Profilometry-based indentation plastometry to obtain stress-strain curves from anisotropic superalloy components made by additive manufacturing. Materialia. 2021;15:101017.
[13] Southern T, Campbell JE, Kourousis KI, et al. Indentation plastometry for study of anisotropy and inhomogeneity in maraging steel produced by laser powder bed fusion. Steel Res Int. 2023;94:2200881.

[14] Campbell JE, Zhang H, Burley M, et al. A critical appraisal of the instrumented indentation technique and profilometry-based inverse finite element method indentation plastometry for obtaining stress–strain curves. Adv Eng Mater. 2021;23:2001496.
[15] Tian Y, Li T, Pang J, et al. Materials design with target-oriented Bayesian optimization. npj Computational Materials. 2025;11:209.
[16] Frazier PI. A tutorial on bayesian optimization, arXiv preprint arXiv:1807.02811 (2018).
[17] Myung JI, Deneault JR, Chang J, et al. Multi-objective Bayesian optimization: a case study in material extrusion. Digit Discov. 2024;4:464–476.
[18] Kiyohara S, Kumagai Y. Bayesian optimization with Gaussian processes assisted by deep learning for material designs. J Phys Chem Lett. 2025;16:5244–5251.
[19] Shahriari B, Swersky K, Wang Z, et al. Taking the human out of the loop: a review of Bayesian optimization. Proc IEEE. 2016;104:148–175.
[20] Lookman T, Balachandran P, Xue D, et al. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. npj Comput Mater. 2019;5:21.
[21] Wang Z, Hutter F, Zoghi M, et al. Recent advances in Bayesian optimization. ACM Comput Surv. 2023;55:1–36.
[22] Khatamsaz D, Vela B, Singh P, et al. Multi-objective materials Bayesian optimization with active learning of design constraints: design of ductile refractory multi-principal-element alloys. Acta Mater. 2022;236:118133.
[23] Pawlak A. Impact of initialization strategies on multi-objective Bayesian optimization in discrete pbf-lb/m process development: a case study on az31 magnesium alloy. Appl Sci. 2025;15(24):12968.
[24] Chernyavsky D, Kononenko DY, Hufenbach JK, et al. Bayesian optimization for laser powder bed fusion of defect-free aa2024. Additive Manufacturing. 2025;114:105022.
[25] Lu C, Jia X, Lee J, et al. Knowledge transfer using Bayesian learning for predicting the process-property relationship of inconel alloys obtained by laser powder bed fusion. Virtual Phys Prototyp. 2022;17:787–805.
[26] Goguelin S, Dhokia V, Flynn JM. Bayesian optimisation of part orientation in additive manufacturing. Int J Computer Integr Manuf. 2021;34:1263–1284.
[27] Erps T, Foshey M, Luković MK, et al. Accelerated discovery of 3d printing materials using data-driven multiobjective optimization. Sci Adv. 2021;7:eabf7435.
[28] Tamura R, Osada T, Minagawa K, et al. Machine learning-driven optimization in powder manufacturing of ni-co based superalloy. Mater Des. 2021;198:109290.
[29] Lee JA, Park J, Sagong MJ, et al. Active learning framework to optimize process parameters for additive-manufactured ti-6al-4v with high strength and ductility. Nat Commun. 2025;16:931.
[30] Zhang R, Song H, Zhang M, et al. Calphad simulation and experimental study on heat treatment process of al-cu-mn-mg-sc cast alloys. Mater Today Commun. 2024;39:108699.
[31] Hu J, Jiang Y, Li J, et al. Alternative acquisition functions of Bayesian optimization in terms of noisy observation. In: *Proceedings of the 2021 European symposium on software engineering, New York, NY, USA*. Association for Computing Machinery, ESSE '21; 2022. p. 112–119.
[32] Naskar S, Suryakumar S, Panigrahi BB. Post-processing of inconel 718 superalloy by laser-based powder bed fusion: microstructures and properties evaluation. Mater Sci Eng A. 2025;921:147601.
[33] Slama C, Abdellaoui M. Precipitation kinetics of and particles in inconel 718 and its influence on mechanical properties. Mater Today Commun. 2024;38:108158.
[34] Le W, Chen Z, Yan K, et al. Early evolution of phase and coarse γ'' phase in inconel 718 alloy with high temperature ageing. Mater Charact. 2021;180:111403.
[35] Gundgire T, Goel S, Klement U, et al. Response of different electron beam melting produced alloy 718 microstructures to thermal post-treatments. Mater Charact. 2020;167:110498.
[36] Hao D, Wang Q, Ma H, et al. The influence of grain boundary character distribution on the high-temperature creep behavior and damage mechanism of inconel 718. Mater Design. 2024;248:113525.
[37] Southern T, Campbell J, Fang C, et al. Use of hardness, pip and tensile testing to obtain stress-strain relationships for metals. Mech Mater. 2023;187:104846.
[38] Bozzolo N, Souaï N, Logé R. Evolution of microstructure and twin density during thermomechanical processing in a γ-γ'' nickel-based superalloy. Acta Mater. 2012;60:5056–5066.
[39] Chlebus E, Gruber K, Kuźnicka B, et al. Effect of heat treatment on the microstructure and mechanical properties of inconel 718 processed by selective laser melting. Mater Sci Eng A. 2015;639:647–655.
[40] Wan H, Yang W, Wang L, et al. Toward qualification of additively manufactured metal parts: tensile and fatigue properties of selective laser melted inconel 718 evaluated using miniature specimens. Journal of Materials Science Technology. 2022;97:239–253.
[41] Anderson M, Thielin AL, Bridier F, et al. Δ phase precipitation in inconel 718 and associated mechanical properties. Mater Sci Eng A. 2017;679:48–55.
[42] De Ath G, Everson RM, Rahat AAM, et al. Greed is good: exploration and exploitation trade-offs in Bayesian optimisation. ACM Trans Evol Learn Optim. 2021;1:22pages.
[43] Rao Z, Bajpai A, Zhang H. Active learning strategies for the design of sustainable alloys. Philos Trans Royal Soc A: Math. Phys Eng Sci. 2024;382:20230242.
[44] Takeno S, Inatsu Y, Karasuyama M. Randomized gaussian process upper confidence bound with tighter bayesian regret bounds (2023).
[45] Bian C, Wang X, Shao W, et al. Adaptive confidence bound based Bayesian optimization via potentially optimal lipschitz conditions. Eng Optim. 2023;55:2051–2069.
[46] Kermanpur A, Babu RP, Larsson H, et al. Simulation and experimental analysis of nanoscale precipitation during ageing treatment of laser powder-bed fusion fabricated in718 ni-based superalloy. Mater Charact. 2022;191:112163.
[47] Cheng Z, Zhou J, Guo G, et al. Machine learning-based investigation of phase precipitation behavior in in718 alloy and its influence on tensile properties. Mater Sci Eng A. 2025;945:148977.
[48] Deng H, Wang L, Liu Y, et al. The evolution law of phase of in718 superalloy in temperature/stress coupled field. Mater Charact. 2022;184:111684.

[49] Popovich V, Borisov E, Popovich A, et al. Impact of heat treatment on mechanical behaviour of inconel 718 processed with tailored microstructure by selective laser melting. Materials Design. 2017;131:12–22.
[50] Huang W, Yang J, Yang H, et al. Heat treatment of inconel 718 produced by selective laser melting: microstructure and mechanical properties. Mater Sci Eng A. 2019;750:98–107.
[51] Fayed E, Shahriari D, Saadati M, et al. Influence of homogenization and solution treatments time on the microstructure and hardness of inconel 718 fabricated by laser powder bed fusion process. Materials. 2020;13:2574.
[52] Franco-Correa JC, Martinez-Franco E, Cruz-Gonzalez CE, et al. Tailored time-temperature-transformation diagram for in718 alloy obtained via powder bed fusion additive manufacturing: phase behavior and precipitation dynamic. Materials. 2023;16:7280.
[53] Trosch T, Strossner J, Volkl R, et al. Microstructure and mechanical properties of selective laser melted inconel 718 compared to forging and casting. Mater Lett. 2016;164:428–431.
[54] Schneider J, Lund B, Fullen M. Effect of heat treatment variations on the mechanical properties of inconel 718 selective laser melted specimens. Addit Manuf. 2018;21:248–254.
[55] Raghavan S, Zhang B, Wang P, et al. Effect of different heat treatments on the microstructure and mechanical properties in selective laser melted inconel 718 alloy. Mater Manuf Processes. 2016;32:1588–1595.
[56] Heo S, Lim Y, Kwak N, et al. Impact of heat treatment and building direction on tensile properties and fracture mechanism of inconel 718 produced by slm process. Metals. 2024;14:440.
[57] Fan Q, Li J, Zheng L, et al. Effect of heat treatment on microstructure and mechanical properties of selective laser melted inconel 718 alloy. PLoS One. 2024;19:e0309156.
[58] Wang X, Li X, Li Z, et al. Superior strength, toughness, and damage-tolerance observed in microlattices of aperiodic unit cells. Small. 2024;20:2307369.
[59] Deng J, Wang X, Zhao R, et al. Golden-ratio–guided aperiodic architected metamaterials with simultaneously enhanced strength and toughness. Adv Funct Mater. 2026;36: e16315.
[60] Balandat M, Karrer B, Jiang DR, et al. Botorch: a framework for efficient monte-carlo Bayesian optimization. Adv Neural Inf Process Syst. 2020;33:21524–21538.
[61] Gardner JR, Pleiss G, Bindel D, et al. Gpytorch: black-box matrix-matrix gaussian process inference with gpu acceleration, 2021.
[62] Gruber K, Ziółkowski GJ, Pawlak AP, et al. Effect of stress relief and inherent strain-based pre-deformation on the geometric accuracy of stator vanes additively manufactured from inconel 718 using laser powder bed fusion. Prec Eng. 2022;76:360–376.
[63] Su J, Chen L, Van Petegem S, et al. Additive manufacturing metallurgy guided machine learning design of versatile alloys. Mater Today. 2025;88:240–250.