ACCSCIENCE PUBLISHING

International Journal of AI for Materials and Design

ORIGINAL RESEARCH ARTICLE

A data-efficient machine learning approach for amorphous Fe-based bulk metallic glass fabrication in powder bed fusion

Jungyeon Kim[1], Sangjun Jeon[1,2], Seong Je Park[3], and Seung Ki Moon[1,4]*

[1]School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore

[2]Department of Future Convergence Engineering, Kongju National University, Cheonan, Republic of Korea

[3]School of Mechanical Engineering, Gyeongsang National University, Jinju, Republic of Korea

[4]Singapore Centre for 3D Printing, School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore

(This article belongs to the *Special Issue: AI Usage in the Analysis of the Additive Manufacturing Process*)

Abstract

The widespread adoption of bulk metallic glasses (BMGs) in aerospace and biomedical industries requires topology-optimized architectures that conventional manufacturing cannot achieve. In response, BMGs have been investigated for powder bed fusion (PBF), but the process remains challenging due to narrow thermal windows, expensive feedstock, and limited data. This study introduces a constrained multi-objective Bayesian optimization framework to optimize key PBF printing parameters, including laser power and scan speed, to maximize hardness while preserving the amorphous state of the printed BMG. Hardness is optimized as the primary objective with density incorporated in the scalarization to regularize the search space, and amorphous retention is enforced through a feasibility probability learned by a logistic classifier. Surrogate models are compared, including Gaussian process, Bayesian additive regression trees, Bayesian multivariate adaptive regression splines (BMARS), and a Bayesian attention neural network. Acquisition scores are computed with constrained expected improvement and are maximized on a uniform grid over power and velocity. Superior predictive accuracy is obtained with BMARS, and 95% credible intervals are calibrated to the measurements. A high-hardness region at high power and low velocity is localized by the surrogates. A fully amorphous sample at 60 W and 1300 mm/s is produced, and a hardness of 1010.4 HV is measured in agreement with the predicted high-hardness band. In conclusion, the study establishes a data-efficient process-window discovery method for BMG PBF, produces an interpretable process map, and demonstrates a screening framework suitable for constrained experimental budgets.

Keywords: Additive manufacturing; Bayesian optimization; Bulk metallic glass; Powder bed fusion

***Corresponding author:**
Seung Ki Moon
(skmoon@ntu.edu.sg)

Citation: Kim J, Jeon S, Park SJ, Moon SK. A data-efficient machine learning approach for amorphous Fe-based bulk metallic glass fabrication in powder bed fusion. *Int J AI Mater Design*. 2025;2(4): 10-23.
doi: 10.36922/IJAMD025390036

Received: September 24, 2025

Revised: October 25, 2025

Accepted: October 30, 2025

Published online: November 13, 2025

1. Introduction

Additive manufacturing allows efficient production of complex parts by directly converting digital designs into solid objects, which allows additive manufacturing to

be widely used in biomedical, automotive, construction, and aerospace industrial applications.[1-4] Among the seven additive manufacturing techniques based on ISO/ASTM 52900, powder bed fusion (PBF) is a printing method suitable for fabricating high-accuracy metal parts.[5,6] PBF achieves high-resolution fabrication by selectively melting fine layers of metal powder using a laser. PBF can fabricate complex geometries such as lattice structures and minimal surfaces that are difficult or impossible to produce with conventional technologies.[7,8] It also improves the usability of various metal materials. Bulk metallic glass (BMG) is an example of a unique material that benefits from PBF. Amorphous alloys of BMG exhibit high strength, hardness, corrosion resistance, and wear resistance.[9-11] However, traditional casting techniques limit the use of BMG to small parts with simple shapes, because BMG must solidify faster than the time it takes for crystals to form.[12-14] PBF provides extremely rapid cooling rates, up to 10^7 K/s, required to retain the amorphous structure, which makes it suitable for producing complex or larger BMG components.[15] However, applying PBF to BMGs remains challenging since high hardness, elastic limit, and corrosion resistance can be achieved only if the melt pool cools at 10^6–10^8 K/s.[16,17] Any local overheating, oxygen uptake, or insufficient enthalpy can trigger nanocrystallization, embrittlement, and cracking.[18,19] Each failed trial is a major loss of time and material and therefore makes data-efficient optimization necessary, since BMG powders are expensive and the amorphous processing window is extremely narrow.

Recent studies have been employing artificial intelligence (AI) approaches to reduce experimental efforts. Supervised convolutional and recurrent networks can classify defects or forecast thermal histories, but the models require thousands of labeled samples.[20-22] Bayesian optimization (BO) overcomes the high data demand by building an uncertainty-aware surrogate and choosing each new experiment where the expected improvement (EI) is highest.[23-26] In projection multiphoton lithography, Johnson *et al.*[27] reduced geometric error to the measurement limit after only four BO iterations by coupling a Gaussian process (GP) surrogate with an EI policy geometry. For laser directed-energy deposition (DED), Karkaria *et al.*[28] embedded a Bayesian long short-term memory surrogate in a digital-twin loop and introduced the Time Series Process Optimization (BOTSPO) algorithm, which achieved optimal ten-parameter laser-power profiles with far fewer function evaluations than Latin-hypercube or Markov chain Monte Carlo (MCMC) searches. Beyond additive manufacturing, Lei *et al.*[24] demonstrated that Bayesian additive regression trees (BART) and Bayesian multivariate adaptive regression splines (BMARS) can outperform GPs when the response surface is high-dimensional or exhibits abrupt transitions, arguing for more flexible surrogates in automated design frameworks. These previous studies confirm that BO delivers the sample efficiency needed for costly experiments and that surrogate choice critically affects convergence.

Despite recent advances, two critical gaps exist in the literature. First, applications of BO to metallic systems have focused almost exclusively on DED or wire-arc additive manufacturing, which leaves PBF unexplored.[29-32] Second, the complex amorphous-crystalline transition with the processability of BMGs has not been modeled in a systematic way.[33-36] Glass formation occurs only when the melt pool cools faster than a critical rate that prevents crystal nucleation. Cooling rates below the critical rate initiate nanocrystals to form and their growth during subsequent layers, which reduces hardness and toughness.[33,37] The resulting discontinuity in the process maps due to nanocrystals limits the development of a smooth surrogate model, but no published study has investigated such a phenomenon. The thermo-kinetic constraints in BMG PBF differ fundamentally from the gradual microstructural trends in Ti-6Al-4V or Inconel 718, so results obtained for conventional alloys cannot be transferred for analysis of BMG. Consequently, PBF still lacks a data-efficient strategy for identifying laser power-speed combinations that preserve glass structure while maximizing mechanical performance.

This study addresses the unmet need by integrating constrained, multi-objective BO with PBF of an FeCrBSiC BMG that has not previously been fabricated using PBF. The proposed framework starts with initial dataset formation through twelve randomly selected laser power and scan speed combinations. Four different surrogate models, which include GP, BART, BMARS, and a Bayesian-Attention Neural Network (BANN), are then trained to predict hardness, while density is included in the scalarization step. A probabilistic classifier is applied to enforce amorphous retention as a feasibility constraint. A constrained expected improvement (CEI) acquisition function is employed to evaluate the balance between hardness and density while excluding regions with high crystallization risk. Iterative retraining of the surrogates after each build is performed to refine the process map and direct the search toward the optimum power and speed region. A final validation build is conducted to confirm the predictive accuracy of the best-performing surrogate.

The four surrogates were chosen to capture smooth trends, sharp transitions, and predictive accuracy under limited data. A GP surrogate was used because GP-based BO showed fast convergence in the printing process and provided calibrated uncertainty for sample-efficient search.

The two non-parametric surrogates, including BART and BMARS, were added to capture local non-linearity and abrupt transitions in the power–speed domain, which were difficult to be well modeled by a smooth GP assumption. A compact BANN was included to provide neural function approximation with posterior uncertainty under a small dataset using approximate Bayesian inference. The selection of the four surrogates aligns with prior BO studies in 3D printing and with adaptive surrogate design in materials experimentation.

This paper is organized as follows: Section 2 discusses the experimental setup and dataset. Section 3 describes the surrogate modeling and optimization procedure. Section 4 reports the results of model comparison and optimization. Section 5 discusses the physical implications and novelty. Section 6 provides the conclusions and future research directions.

2. Data collection and methodology

2.1. Experimental setup and conditions

The powder and equipment employed were FeCrBSiC (Dura-Metal, Singapore) and EOS M100 (EOS GmbH, Germany), respectively. Laser energy densities were experimentally selected between 25.2 J/mm^3 and 35.7 J/mm^3 for sample fabrication, which were calculated using the volumetric energy density (E) equation:

$$E = \frac{P}{vht} \quad \text{(I)}$$

Where P represents the laser power, v is the scan speed, h is the hatch distance, and $t$is the layer thickness. E represents the average energy delivered per unit volume of material during one layer printing and serves as an integrated measure of process input. The hatch distance and layer thickness were fixed at 70 μm and 20 μm, respectively. Laser energy densities below 25.2 J/mm^3 were found to be insufficient to achieve full melting, leading to poor layer bonding and incomplete builds, which lowered both hardness and surface density. However, E values above 35.7 J/mm^3 resulted in excessive energy input, causing defects such as cracking due to thermal stress accumulation and promoting local crystallization. Crystallization led to embrittlement and a decrease in attainable hardness. Besides, three representative laser powers (60 W, 70 W, and 120 W) were randomly selected to explore the effects of laser power and laser scan speed on the resulting part quality. At each laser power, the scan speed was varied within the defined volumetric energy density range to fabricate 10 × 10 × 10 mm^3 cubic samples. After fabrication, the samples were separated from the substrate using wire cutting. X-ray diffraction (XRD) analysis was performed using a SmartLab diffractometer (Rigaku, Japan) with Cu Kα radiation (λ = 0.1541 nm). The XRD patterns were recorded in the 2θ range of 25°–90°. In contrast, for density and hardness measurements, the samples were mounted and polished to ensure smooth cross-sectional surfaces. The surface densities were evaluated by analyzing cross-sectional images obtained via an optical microscope, Olympus DP22 (Evident, Japan), using ImageJ software. Vickers hardness tests were performed on the same polished surfaces via a Microscan hardness tester (Omron, Japan), respectively. The resulting dataset served as the experimental foundation for the machine learning-based optimization process described in the subsequent sections.

2.2. Surrogate model construction and data pre-processing

In this study, a Bayesian surrogate modeling approach is employed to capture the relationships between PBF process parameters and the resulting material responses, including hardness, density, and amorphous retention. A surrogate that can learn from sparse input data and quantify predictive uncertainty is required because the high cost and lengthy duration of BMG PBF trials limit the number of feasible experiments. Therefore, GP regression is chosen as a primary surrogate model. The GP surrogate models the unknown objective function in a Bayesian manner, where GP assumes any set of function values has a joint Gaussian distribution defined by a mean function and a covariance kernel.[29] After observing data, the GP posterior provides a predictive mean μ and variance σ^2 for any candidate input x, which serve as the model's prediction and uncertainty estimate, respectively. An automatic relevance determination radial basis function kernel with a constant and white-noise term was used for GP. The kernel parameters were optimized by maximizing log-marginal likelihood with 15 random restarts. Uncertainty quantification is crucial in BO, since predictions with the lowest uncertainty are used inside the algorithm to identify the most favorable predicted outcome, and predictions with high uncertainties provide meaningful information for new data.

While GPs are suitable for smooth, continuous functions, the BMG PBF can exhibit non-stationary or discontinuous behaviors such as sudden changes in properties near processing limits, and heteroscedastic noise. A standard GP with a stationary kernel often fails to capture non-stationarity and input-dependent noise.[24,38,39] Therefore, more flexible nonparametric Bayesian surrogates are evaluated to address the non-stationarity and heteroscedastic noise: BMARS and BART. BMARS and BART can capture complex nonlinear relationships and

interactions that a simple GP cannot represent. BMARS constructs a flexible polynomial basis to model localized trends, whereas BART integrates many shallow trees to accommodate sudden changes and complex interactions. Both BMARS and BART perform automatic variable selection and include Bayesian regularization, which helps avoid overfitting even with limited data. Both BMARS and BART implement a fully Bayesian inference procedure, which produces a predictive distribution at each input. BMARS employed 40 maximum terms, a second-degree basis, and a penalty of 0.5. BART used 30 trees with 60–200 burn-in and posterior draws. The resulting uncertainty quantification aligns with the GP surrogate's mean-variance output and supports direct integration into the acquisition computations.

Besides, BANN is also studied, which is an additional nonparametric model alongside GP, BMARS, and BART. A compact BANN surrogate is developed, which is designed specifically for the limited dataset of this study. One hidden layer with a rectified linear unit (ReLU) is used. ReLU is chosen to avoid saturation and to keep a locally linear function class that controls capacity under small datasets, which is consistent with practical guidance for small-data BANN.[40] Dropout of 0.25 is kept active during inference to approximate Bayesian inference and to provide calibrated epistemic uncertainty.[41] Two output heads are used to predict hardness and surface density, and each head produced a mean $\mu(x)$ and a log variance $s(x)$ where x is defined as the two-dimensional input process vector (P, v). Training is performed with the Gaussian negative log likelihood with L2 weight, which is expressed as follows:[41,42]

$$\mathcal{L}(x,y)=\frac{1}{2}s(x)+\frac{1}{2}\frac{\left(y-\mu(x)\right)^2}{\exp\{s(x)\}}+\lambda\theta_2^2 \quad \text{(II)}$$

In Equation II, y represents the observed target, hardness, and surface density. λ is a weight-decay coefficient that sets the strength of L2 regularization, θ represents the set of all trainable weights in the neural network, and $\|\theta\|_2^2$ represents the squared L2 norm or sum of squares of all weights in θ. λ is set to 0.005, learning rate is 0.003, and the early stopping patience is fixed at 60 epochs.

At prediction, M stochastic forward passes with dropout were used. The predictive mean $\bar{\mu}(x)$ is expressed as follows:

$$\bar{\mu}(x)=\frac{1}{M}\sum_{m=1}^{M}\mu^{(m)}(x) \quad \text{(III)}$$

Where $\mu^{(m)}(x)$ is the predicted mean from the m-th stochastic pass with dropout, and M is the number of Monte Carlo passes.[40,41] This estimator was used as the MC-dropout approximation to Bayesian inference.[40,41]

Separate surrogate models predict hardness and density. Input-dependent noise is incorporated by allowing predictive uncertainty to vary across the process parameter space rather than imposing constant observation noise. A heteroscedastic GP formulation models the noise variance as a function of the input, for example, via an auxiliary GP over the noise term. Experimental observations indicate that regions of very high laser power or scanning speed yield greater outcome variability from unstable melt pool dynamics, whereas other regions behave more consistently. Local adaptation of the noise model improves uncertainty quantification because predictive variance rises where data is sparse or noisy and falls in well-characterized regions. The heteroscedastic modeling proves to be critical for BMG printing, where abrupt shifts in process conditions can drastically affect part quality variability.

Amorphous phase preservation is calculated as a probabilistic feasibility constraint in the optimization process rather than as a distinct objective. An ideal printer parameter combination must preserve the minimum critical amorphous fraction in the printed part. A probabilistic feasibility surrogate predicts whether a given set of process parameters produces an amorphous structure. Each experiment is labeled with a binary indicator , which describes the structural order between 0 and 1 as follows:

$$\Delta=\begin{cases}0\ (crystalline)\\ 1\ (amorphous)\end{cases} \quad \text{(IV)}$$

A logistic classifier learns a latent function $g(x)$ from the feasibility labels. The feasibility probability at x equals $\Phi[g(x)]$, with the standard-normal cumulative distribution evaluated at $g(x)$. The predicted probability $p(x)$ gives a continuous estimate from zero to one of meeting the amorphous preservation requirement. Training on feasible versus infeasible outcomes produces predictive uncertainty in regions with limited or conflicting data. The logistic classifier delivers both a feasibility prediction and a confidence level. Constrained BO multiplies the EI metric by $p(x)$ so that parameter settings predicted to violate the amorphous criterion receive a negligible acquisition score. Modeling the feasibility constraint via a probabilistic classifier follows a standard practice in constrained BO, and the acquisition function balances potential objective gains against the requirement to preserve the amorphous phase while accounting for uncertainty.

Four surrogate models, including GP, BART, BMARS, and BANN, are developed to model the physical responses under limited data and to provide predictive uncertainty

for BO. Data pre-processing is applied before model fitting. Raw measurements of laser power, scan speed, Vickers hardness, and surface density are used as recorded. GP targets are normalized using the model's internal normalization option. Vickers hardness and surface density were standardized for training and were inverse-scaled at prediction for BART and BANN. Targets were kept in original units for BMARS. The feasibility classifier for amorphous retention is trained on original power and scan speed with binary labels.

2.3. Acquisition and optimization procedure

Figure 1 illustrates the overall sequential workflow, which collects process data; trains surrogate models, applies Pareto efficient global optimization (ParEGO) scalarization, evaluates CEI, and selects the next build parameters. Surrogate models for hardness, density, and amorphous-phase feasibility are supplied to a constrained multi-objective BO that identifies optimal PBF process parameters. The optimization strategy balances two requirements: (i) concurrent improvement of hardness and density, and (ii) compliance with the amorphous-phase feasibility requirement. The ParEGO method is employed to reduce multiple objectives to a single-objective acquisition problem at each iteration.[42] Besides, scalarization is used to combine the surrogate predictions into a weighted sum, where a weight vector,, balances hardness and density. At each iteration, a random weight between 0 and 1 is drawn from the uniform distribution, and the composite objective is defined as:

$$f_{scalar}(x) = w \cdot \tilde{H}(x) + (1-w) \cdot \tilde{D}(x) \quad \text{(V)}$$

Where $\tilde{H}(x)$ and $\tilde{D}(x)$ are the predicted hardness and density, respectively. Variation of the weight w across iterations explores Pareto trade-offs between the two metrics: a larger w places more emphasis on predicted hardness at the expense of density, whereas a smaller w places more emphasis on predicted density at the expense of hardness. After multiple iterations with distinct random weights, ParEGO produces a solution set that approximates the Pareto-optimal frontier for both objectives. ParEGO extends single-objective EI using the same GP surrogates and acquisition framework and proves effective under the limited experimental feasibility, which is typical for BMG PBF research.

After scalarization, the multi-objective problem becomes a single-objective constrained optimization at each iteration, with the selected weight defining the hardness-density trade-off. Selection of the next experimental condition proceeds via maximization of the acquisition function, which evaluates each candidate according to the surrogate's predicted mean $\mu(x)$ and variance $\sigma^2(x)$. The EI criterion captures both exploitation and exploration by computing the expected gain over the current best value f^*. For a candidate x, the improvement random variable is defined as:

$$I(x) = \max\left(f_{scalar}(x) - f^*, 0\right) \quad \text{(VI)}$$

Where $f_{scalar}(x)$ represents the surrogate's predictive distribution. Under a Gaussian posterior, the $EI(x)$ is described in a closed form as:

$$EI(x) = \left(\mu(x) - f^*\right)\Phi(Z) + \sigma(x)\phi(Z) \quad \text{(VII)}$$

Where $\sigma(x)$ is the predicted standard deviation, and Z is the measure of the distance between the surrogate's predicted mean and f^*.[43] The functions Φ and ϕ represent the cumulative distribution function and probability density function of a standard normal distribution, respectively. The closed-form results from integrating the improvement variable under the surrogate's Gaussian posterior. The first term, $(\mu(x)\text{-}f^*)\,\Phi(Z)$, captures the expected gain if the predicted mean exceeds the current best, weighted by the corresponding probability. The

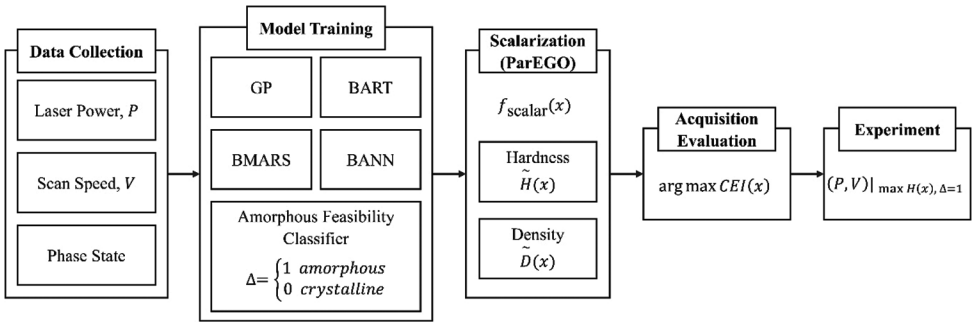

Figure 1. Constrained multi-objective Bayesian optimization workflow for BMG PBF parameter selection
Abbreviations: BANN: Bayesian attention neural network; BART: Bayesian additive regression trees; BMARS: Bayesian multivariate adaptive regression splines; BMG: Bulk metallic glass; CEI: Constrained expected improvement; GP: Gaussian process; ParEGO: Pareto efficient global optimization; PBF: Powder bed fusion.

second term, $\sigma(x)\phi(Z)$, represents the value of sampling in regions with high uncertainty, where unexpectedly large improvements remain possible. Consequently, $EI(x)$ attains large values in regions of high predicted performance or high uncertainty, which provides a balance between exploitation and exploration.

A feasibility constraint on amorphous-phase preservation multiplies the EI criterion. The logistic classifier estimates $p(x)$, which is the probability that parameter vector x yields an amorphous build. The CEI can be computed as:

$$CEI(x) = p(x)\,EI(x) \quad \text{(VIII)}$$

A candidate x with a large EI but low feasibility probability receives a near-zero CEI score. Only parameter settings with high $p(x)$ can retain the full EI value. As a result, the acquisition function balances gains in hardness and density against the requirement to preserve the amorphous structure. Regions where $p(x)$ is close to 1 reduce CEI to standard EI, whereas regions where $p(x)$ is close to 0 yield negligible CEI regardless of predicted performance.

Maximization of the CEI employs a uniform grid over the laser-power and scan speed ranges. The acquisition score $CEI(x)$ is computed at each grid point in milliseconds, and the point with the maximum score becomes the next experimental condition. A full-grid search identifies the maximum global CEI within the grid resolution and avoids the risk of settling in local maxima that can affect gradient-based or stochastic search methods.

2.4. Statistical analysis

All computations ran on a Windows 10 laptop equipped with an Intel Core i5 13420H (2.10 GHz) and 16 GB RAM. The Python 3.6 environment used scikit-learn 1.3 for Gaussian-process regression and logistic regression, bartpy for BART, pyEarth 0.1 for BMARS, and TensorFlow Probability 0.22.5 for the BANN. A global NumPy random seed fixed every stochastic step in the workflow.

Model accuracy was evaluated using leave-one-out cross-validation (LOOCV) for reliable estimation of predictive performance under the small sample condition. LOOCV provides reliable Bayesian predictive accuracy relative to information-criteria approaches for limited data.[44,45] In each iteration, one observation was used as the test case, and the remaining observations served as the training set. The procedure was repeated over all observations. The average prediction error across iterations provided a stable estimate of generalization.

Root mean square error (RMSE) and coefficient of determination (R^2) were computed from the LOOCV results for both hardness and surface density predictions. The logistic classifier that enforces the amorphous-phase constraint was also validated under the same LOOCV framework.

BO progress was tracked by recording the best hardness observed after each iteration. Matplotlib produces all figures, and pandas records the accuracy and optimization tables. No further hypothesis testing, calibration scoring, or predictive interval assessment is necessary for the objectives of this study.

3. Results

3.1. Model comparison

Figure 2 illustrates the predicted hardness response surfaces as functions of laser power and scan velocity for each surrogate model: GP, BMARS, BANN, and BART. The figure presents four distinct surfaces, each representing a different modeling approach applied to the same initial dataset, with warmer colors indicating higher hardness values.

Figure 2A displays the GP surrogate model results. The GP surface exhibited a smooth gradient, with hardness values rising from approximately 700 HV at low power and high velocity to over 1000 HV at high power and low velocity. Hardness increased with laser power and decreased with scan velocity in the GP predictions, which is consistent with the stationary kernel's ability to represent the dominant process-property relationship. The smoothness of the GP surface reflected its underlying assumption of continuous, gradual changes in the response, which proved suitable for the observed data. Figure 2B shows the BMARS surrogate model, which showed a similar overall gradient to the GP but with sharper transitions between regions of differing hardness. Hardness increased significantly with higher power and lower velocity, where the value reached its maximum in the same region. The spline-based formulation of BMARS allowed the model to fit both gradual and abrupt changes, resulting in a surface that captured the main trend while also highlighting localized variations. The flexibility of BMARS produced a clearer separation between high- and low-hardness regions compared with the GP model.

BANN surrogate model predictions are shown in Figure 2C. The BANN surface appeared as a broad plateau across the domain with multiple local peaks and valleys. A gentle rise was found toward high power and low velocity regions, but the model did not form a clear ridge. The response range was compressed, indicating limited sensitivity to the inputs compared with GP and BMARS. Figure 2D illustrates the BART surrogate model,

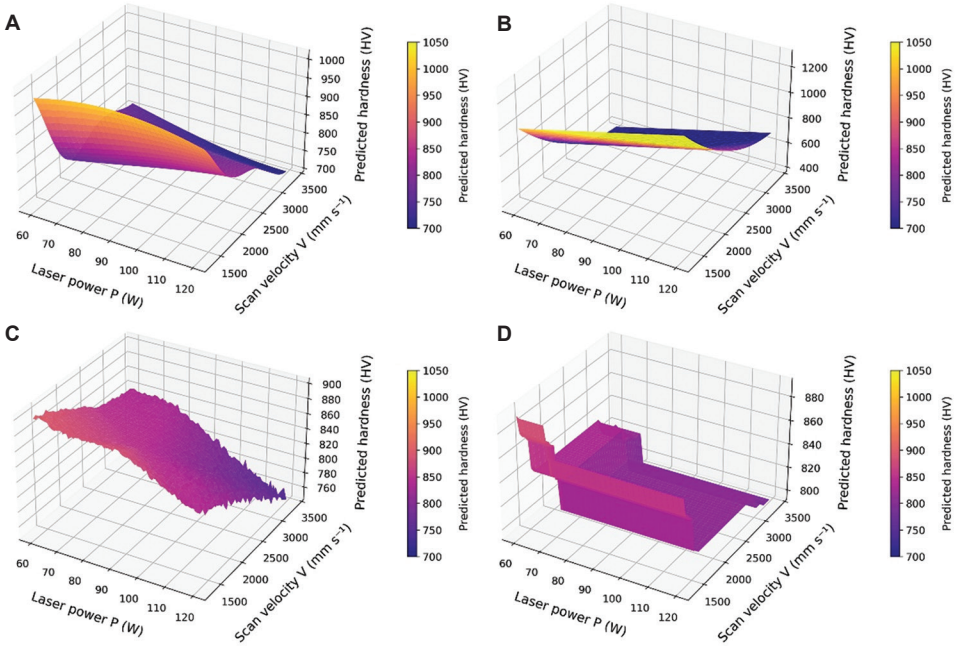

Figure 2. Predicted hardness response surfaces as functions of laser power (P) and scan velocity (V) for each surrogate model: (A) Gaussian process, (B) Bayesian multivariate adaptive regression splines, (C) Bayesian attention neural network, and (D) Bayesian additive regression trees

which partitions the design space into piecewise-constant regions of uniform hardness. A distinct high-hardness surface appeared at high power and low velocity, sharply separated from a lower-hardness regime elsewhere. The tree-ensemble structure of BART displayed abrupt changes in hardness and facilitated the identification of discrete processing regimes. The stepwise approximation effectively captured sharp transitions in the response surface.

All four models successfully identified the peak-hardness region at high power and low velocity, but the models differed in representation and smoothness. The GP and BMARS surrogates provided smooth interpolations of the underlying trend, where BMARS showed slightly sharper regional transitions. The BART model outperformed in isolating distinct hardness regimes through its stepwise approximation, whereas the BANN model exhibited excessive variability across the domain.

3.2. Surrogate-model accuracy

Figure 3 shows predicted hardness against laser power at 1500 mm/s. Blue curves are posterior means. Shaded regions mark the 95% credible intervals. Red symbols are the measured specimens.

The BMARS captured the monotonic rise in hardness as illustrated in Figure 3B. Hardness increased from about 880 HV at 60 W to roughly 1230 HV at 120 W, which matched the experimental data. The 95% confidence interval remained narrow, and every experimental data point lay inside the band. The tight span indicates precise uncertainty quantification. The GP model in Figure 3A followed the same trend but appears smoother. The GP model underestimated the peak by nearly 50 HV, with a broader confidence interval. The wider band indicated higher epistemic uncertainty.

The BANN model in Figure 3C showed weak dependence on power. The posterior mean remained near a flat level across the range. The 95% confidence interval was wide and enclosed the two high-hardness data points. The model did not recover the upward trend and reflected low confidence in the prediction rather than overconfidence. Besides, the BART model in Figure 3D predicted a flat segment beyond 90 W. The mean deviated from the highest experimental data point by more than 100 HV. The BART model's interval stayed narrow but excluded the data points. The model understated both the mean and the variance in the high-power regime.

The prediction performances of four different models are summarized in Table 1. BMARS yielded an RMSE of 36 HV and an R^2 of 0.89. The GP followed with RMSE of 46 HV and R^2 of 0.83. BART and BANN exceeded 90 HV in RMSE and fell below 0.35 in R^2 values.

3.3. Optimization outcomes

Figure 4 compares the measured and predicted hardness values for each surrogate under LOOCV. Points distributed near the dashed 45° lines represent accurate

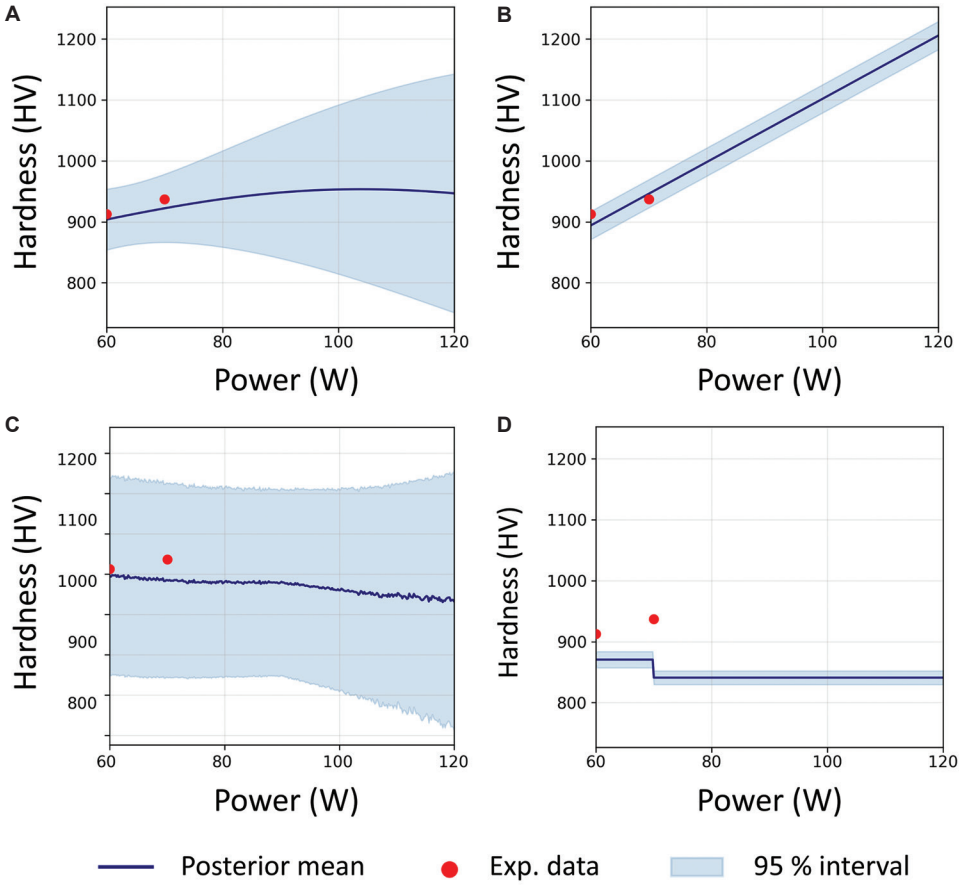

Figure 3. Hardness predictions of different models at scan velocity of 1500 mm/s: (A) Gaussian process, (B) Bayesian multivariate adaptive regression splines, (C) Bayesian attention neural network, and (D) Bayesian additive regression trees

Table 1. Surrogate accuracy metrics

Surrogate	RMSE (HV)	R^2
GP	45.59	0.832
BMARS	36.37	0.893
BANN	92.77	0.303
BART	96.59	0.244

Abbreviations: BANN: Bayesian attention neural network; BART: Bayesian additive regression trees; BMARS: Bayesian multivariate adaptive regression splines; GP: Gaussian process; RMSE: Root mean square error.

predictions. Both GP and BMARS models followed the ideal trend closely, whereas BANN and BART showed larger deviations. The GP points concentrated around the diagonal but slightly underestimated the highest hardness regions, whereas BMARS aligned most tightly with the measurements. The BANN model compressed the range of predicted values, and the BART model produced wider scatter that did not capture the full variation of the data. The plots confirm that BMARS and GP achieved the best agreement with experimental hardness, which is also consistent with the quantitative ranking in Table 1.

Figure 5 illustrates the running best hardness that each BO model recorded after every evaluation. BANN reached its final level by the third iteration at about 1080 HV. BMARS and GP were also stabilized by iteration 1 near 1027 HV and 1017 HV. BART reached its final level by iteration 2 near 1053 HV. Traces remained flat afterward, which indicates early convergence at different plateaus.

The rapid initial increase in the running best hardness during the first two to four iterations was consistent with the shapes of the surrogate response surfaces, as shown in Figure 2. A steep gradient near the high power and low velocity region was learned by BANN, and the region was prioritized by the acquisition. The same zone was reached by BART, but a lower peak was produced because the stepwise approximation underestimated the maximum. Curvature was preserved by BMARS, and the peak along the high hardness ridge was smoothed. The GP surface was smoothed more strongly, and the inflection near the crystalline boundary was blunted. Figures 2 and 3 show the same pattern for early entry into the high power and low velocity region and for peak height. GP and BMARS reached the region first, where BMARS captured sharper

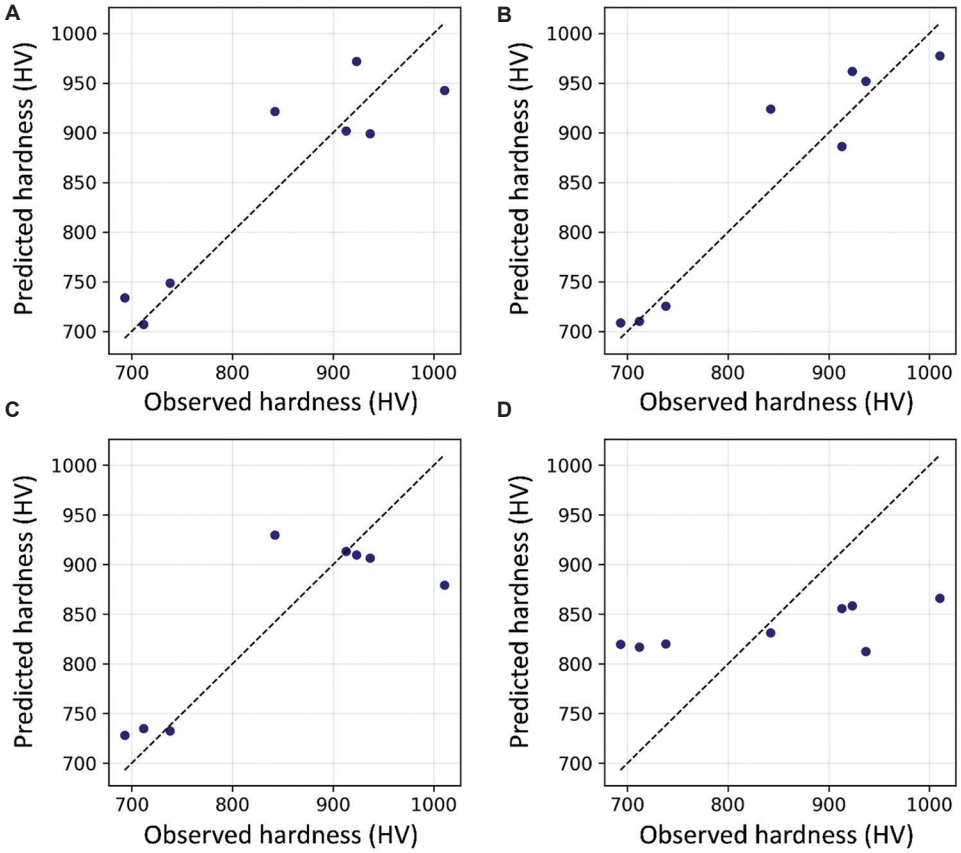

Figure 4. Observed versus predicted hardness for the four surrogate models under leave-one-out cross-validation: (A) Gaussian process, (B) Bayesian multivariate adaptive regression splines, (C) Bayesian attention neural network, and (D) Bayesian additive regression trees

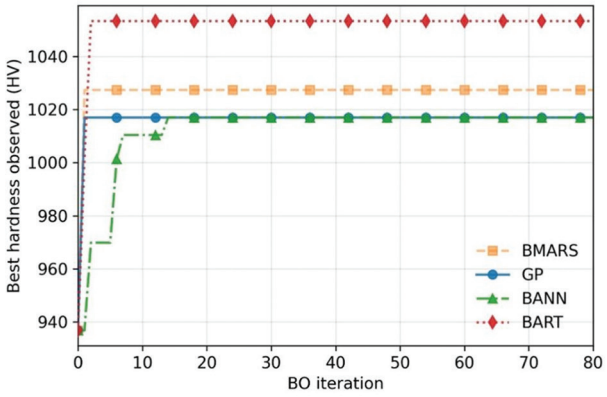

Figure 5. Optimization convergence trajectories for each surrogate model Abbreviations: BANN: Bayesian attention neural network; BART: Bayesian additive regression trees; BMARS: Bayesian multivariate adaptive regression splines; BO: Bayesian optimization; GP: Gaussian process.

local changes, GP smoothed the peak, BART reached the highest hardness peak, and BANN reached the optimum with the highest number of iterations.

Table 2 lists four quantities for each model, including maximum hardness, the first iteration at or above 95% of the maximum, and the recommended power and velocity. GP and BMARS reached the 95% criterion by iteration 1. BART met the criterion by iteration 2, and BANN reached the criterion by iteration 3.

Table 2. Optimization results

Surrogate	Maximum hardness (HV)	Iteration	Recommended power (W)	Recommended velocity (mm/s)
GP	1017	1	64	1300
BMARS	1027.4	1	66	1300
BANN	1017	3	60	1350
BART	1053.3	2	61	1300

Abbreviations: BANN: Bayesian attention neural network; BART: Bayesian additive regression trees; BMARS: Bayesian multivariate adaptive regression splines; GP: Gaussian process.

Near the crystalline boundary, a low probability of amorphous retention was predicted by the logistic classifier. EI was multiplied by the predicted probability, and the constrained score was near zero for power and velocity settings located in the boundary zone. Power and velocity settings in the boundary zone were ranked poorly

and were not selected. Sampling was concentrated along the ridge where high predicted hardness coincided with high feasibility. Sharp transitions in the BMARS and BART surfaces were mirrored in the selected evaluations.

A validation build was produced at a laser power of 60 W, a scan velocity of 1300 mm/s, and a hardness of 1010.4 HV was measured. A Vickers indentation from the build is shown in Figure 6. The value lay on the high hardness band in Figure 7 and confirmed movement toward a physically meaningful region.

Figure 7 maps predicted hardness across power and velocity using BMARS. Hardness values above 950 HV occupied the high-power and low-velocity region and appear as an elongated high-hardness band. Iso-hardness contours were oriented diagonally from higher power toward lower velocity and were nearly parallel to lines of constant volumetric energy density from Equation I. Hardness decreased when power fell below 80 W velocity exceeded 2600 mm/s. Beyond a volumetric density of about 60 J/mm^3, a saturation response was observed, and comparable hardness values were produced by many power and velocity settings. Trade-offs were present between hardness, build rate, and energy demand because regions with higher hardness often required lower velocity and higher power, which reduced throughput and increased energy input.

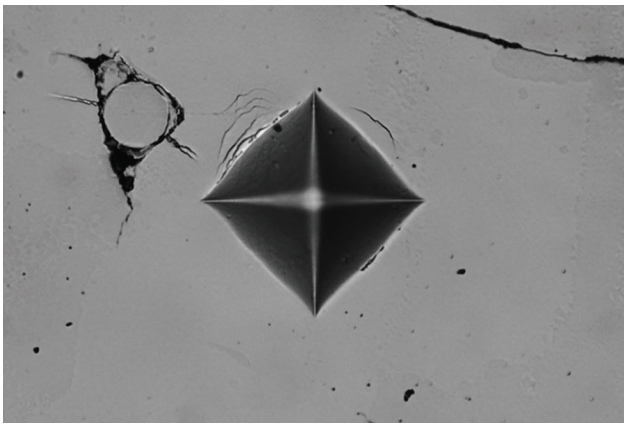

Figure 6. Vickers hardness indentation mark on the sample fabricated with a laser power of 60 W and a scanning speed of 1300 mm/s

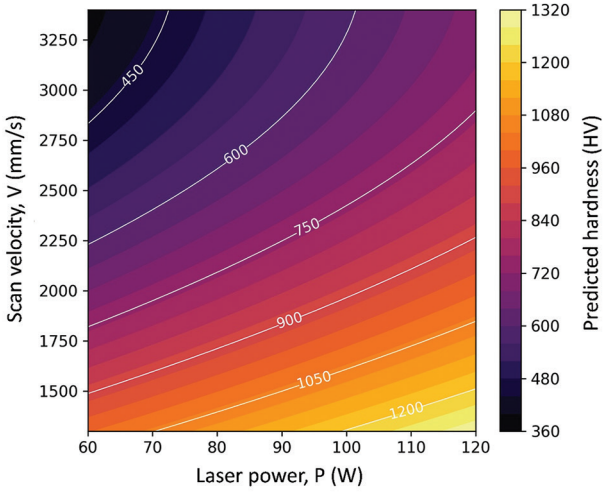

Figure 7. BMARS-predicted hardness across laser power and scan velocity
Abbreviation: BMARS: Bayesian multivariate adaptive regression splines.

4. Discussion

BART completed the search the fastest, and BMARS, GP, and BANN followed in that order. Although BMARS iteration finished second, the model produced the smallest error and the most reliable confidence intervals. The flexible spline bases in BMARS captured the broad rise in hardness as energy increased and preserved the sharp inflection point where crystallization began. In contrast, neither the stationary GP kernel nor BANN captured the sharp crystallization transition, as the models provided a smooth surface, a near-flat surface, respectively. Similar behavior appeared in high-dimensional MAX-phase screening, where Lei *et al.*[24] reported that BMARS and BART used fewer than 0.3% of the design space, whereas GP variants often ran out of budget BOTSPO. This research confirms that adaptive knots and tree splits give the surrogate the local freedom required when responses change abruptly.

The optimization results show that BMARS guided the search toward conditions that combined high predicted hardness with a high probability of amorphous retention. This focus avoided sampling power-velocity settings that thermokinetic theory identifies as infeasible and accelerated progress toward the best-performing region. Similar behavior has been reported in other additive manufacturing applications, such as projection multi-photon lithography and the BOTSPO digital-twin framework for DED, where BO reduced the number of required trials compared with conventional parameter sweeps.[28] In the present study, the method reached a setting that produced approximately 1030 HV hardness in physical builds while keeping the experimental effort low, demonstrating its potential for efficient process development in BMG PBF.

Hardness increased with volumetric energy density until about 60 J/mm^3 and remained nearly constant beyond that point. Above this volumetric energy density threshold, the melt pool temperature increased while the cooling rate stayed nearly unchanged, so the glass-forming driving force did not increase further. Below 80 W or

above 26 mm/s, hardness and density decreased because local cooling fell below the critical rate and nanocrystals formed. The logistic interaction term extracted from the classifier showed that raising scan speed compensates part of the excess power by shortening solidification time, which matches continuous-cooling-transformation theory for Fe-based glasses.

The optimum laser power and scan speed combination was located at the highest volumetric energy density input that maintained glass formation, and it corresponded to the first high-power pulse in the time-dependent power profiles selected by BOTSPO. Both static BMARS map and dynamic BOTSPO profiles converge on the same physical principle, which is to apply the highest energy input to close pores, and then cool rapidly to preserve the amorphous state. The consistency between the BMARS map and the BOTSPO profiles indicates that the BMARS map can provide a basis for initializing a process digital twin for PBF, which can then be refined layer by layer with *in situ* monitoring data, as demonstrated for DED control using the BOTSPO framework.[28]

The novelty of this study is defined by three main contributions. It is the first study to apply BO to PBF of BMG, which is a material system with an extremely narrow amorphous window for processability. The proposed framework introduces a constrained multi-objective approach that optimizes hardness and density while enforcing amorphous retention as a probabilistic feasibility condition. Furthermore, four surrogate models are systematically compared in terms of predictive accuracy, uncertainty calibration, and optimization efficiency, which establishes how surrogate selection influences convergence under limited experimental data. Overall, the contributions outlined above expand constrained BO's applicability to metastable alloys like BMG, which provides a data-efficient approach for process design in additive manufacturing under experimental constraints.

5. Conclusion

The objective of this study was to identify process parameters that preserved the amorphous structure of FeCrBSiC during PBF. The proposed approach was applied to a constrained multi-objective BO with surrogate modeling. This study established a data-efficient strategy for printing fully amorphous FeCrBSiC by PBF. A constrained multi-objective BO loop was applied to 80 PBF trials and identified a process condition that retained an amorphous FeCrBSiC structure with the maximum surface density and produced roughly 1030 HV Vickers hardness. Among the four surrogate models, BMARS exhibited the lowest RMSE, the highest R^2, and the only 95 % credible band that enclosed every measured value, which confirmed the superior fidelity compared to other models.

The experimental validation at 60 watts and 1300 mm/s further highlighted the predictive strength of the spline surrogate. The resulting contour map defined a narrow ridge of high hardness along the high-power, low-velocity border, which gave a clear processing window for engineering practice. A probabilistic classifier for amorphous-phase feasibility is combined with an uncertainty-aware acquisition function to determine experimental parameter selection. The combination of probabilistic feasibility modeling and uncertainty-aware acquisition makes the methodology applicable to other metastable alloys that require rapid solidification and precise energy control, and also supports the evaluation of new alloy systems for PBF processing.

Several limitations remain in this research. Hatch spacing and layer thickness were constant in this research, which makes keyhole porosity and track-shape effects unexplored. The surrogate was trained on one type of alloy, but transfer learning could reuse spline hyper-priors for alloys of lower glass-forming ability. Finally, hardness alone cannot guarantee structural reliability, so future work should also quantify fracture toughness and corrosion along the same energy range.

Future research will replace the squared exponential kernel with a spectral mixture or Gibbs kernel so that the length scale varies across the power and velocity domains. Reversible jump Markov chain Monte Carlo (RJMCMC) will detect change points that mark the transition from glass formation to crystallization and will switch kernel hyperparameters on each side. A hierarchical prediction derived from melt pool heat flow calculations will regularize predictions in regions with sparse data and will introduce parameters that correspond to physical quantities such as cooling rate. A multi-output GP will integrate hardness, density, and amorphous fraction, which will preserve coherent uncertainty when the optimizer chooses the next build. The acquisition function will account for estimated machine time and powder mass and will schedule batches that use fewer resources while still guiding the search toward the Pareto front. An embedded multi-sensor fusion-based pipeline will read optical and acoustic signals throughout the build, update the surrogate after each layer through Bayesian filtering, and adjust laser power and scan speed in real time to keep the melt pool within the glass-forming domain whenever drift emerges.

Acknowledgments

This research is supported by the Singapore Centre for 3D Printing (SC3DP) and the Ministry of Trade, Industry

and Energy (MOTIE, Korea) through the 2025 R&D HRD Program for Problem-Solving Based on Future Mobility Innovation Technology Development (Project No. P0023726).

Funding

This research is conducted by the Industrial Technology Innovation Program (KEIT project no. 20024344, Development of AI-based high carbon steel alloy design and sintering-based additive manufacturing technology for 7.0 L/Hr-level high-speed production of powertrain components with tensile strength over 1.0 GPa in the next-generation mobility) funded by the Ministry of Trade, Industry & Energy of the Republic of Korea.

Conflict of interest

Seung Ki Moon is an Editorial Board Member of this journal, but was not in any way involved in the editorial and peer-review process conducted for this paper, directly or indirectly. Separately, other authors declared that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Author contributions

Conceptualization: Jungyeon Kim
Formal analysis: Jungyeon Kim, Sangjun Jeon
Funding acquisition: Seung Ki Moon
Investigation: Jungyeon Kim, Seong Je Park
Methodology: Jungyeon Kim, Sangjun Jeon, Seong Je Park
Software: Jungyeon Kim
Writing–original draft: Jungyeon Kim, Sangjun Jeon, Seong Je Park
Writing–review & editing: Seong Je Park, Seung Ki Moon

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data

The code and datasets supporting the findings of this study are available from the corresponding author upon reasonable request.

References

1. Tofail SA, Koumoulos EP, Bandyopadhyay A, Bose S, O'Donoghue L, Charitidis C. Additive manufacturing: Scientific and technological challenges, market uptake and opportunities. *Mat Today*. 2018;21(1):22-37.

   doi: 10.1016/j.mattod.2017.07.001

2. Kumar R, Kumar M, Chohan JS. The role of additive manufacturing for biomedical applications: A critical review. *J Manuf Process*. 2021;64:828-850.

   doi: 10.1016/j.jmapro.2021.02.022

3. Blakey-Milner B, Gradl P, Snedden G, *et al*. Metal additive manufacturing in aerospace: A review. *Mater Des*. 2021;209:110008.

   doi: 10.1016/j.matdes.2021.110008

4. Rehman M, Yanen W, Mushtaq RT, *et al*. Additive manufacturing for biomedical applications: A review on classification, energy consumption, and its appreciable role since COVID-19 pandemic. *Prog Addit Manuf*. 2023;8(5):1007-1041.

   doi: 10.1007/s40964-022-00373-9

5. *ISO/ASTM 52900:2021(en), Additive Manufacturing - General Principles — Fundamentals and vocabulary*. Available from: https://www.iso.org/obp/ui/#iso: std:iso.astm:52900:ed-2:v1:en [Last accessed on 2025 Aug 10].

6. Wang J, Jeong SG, Kim ES, Kim HS, Lee BJ. Material-agnostic machine learning approach enables high relative density in powder bed fusion products. *Nat Commun*. 2023;14(1):6557.

   doi: 10.1038/s41467-023-42319-x

7. Park SJ, Heogh W, Yang J, *et al*. Meta-structure of amorphous-inspired 65.1Co28.2Cr5.3Mo lattices augmented by artificial intelligence. *Adv Compos Hybrid Mater*. 2024;7(6):1-22.

   doi: 10.1007/s42114-024-01039-6

8. Park SJ, Lee JH, Yang J, *et al*. Lightweight injection mold using additively manufactured Ti-6Al-4V lattice structures. *J. Manuf Process*. 2022;79:759-766.

   doi: 10.1016/j.jmapro.2022.05.022

9. Howard J, Carlson K, Chidambaram D. High-temperature metallic glasses: Status, needs, and opportunities. *Phys Rev Mater*. 2021;5(4):040301.

   doi: 10.1103/physrevmaterials.5.040301

10. Scully JR, Gebert A, Payer JH. Corrosion and related mechanical properties of bulk metallic glasses. *J Mater Res*. 2007;22(2):302-313.

    doi: 10.1557/jmr.2007.0051

11. Xu T, Pang S, Li H, Zhang T. Corrosion resistant Cr-based bulk metallic glasses with high strength and hardness. *J Non Cryst Solids*. 2015;410:20-25.

    doi: 10.1016/j.jnoncrysol.2014.12.006

12. Chen M. A brief overview of bulk metallic glasses. *NPG Asia Mater*. 2011;3(9):82-90.

    doi: 10.1038/asiamat.2011.30

13. Mi XL, Hu L, Wan ZX, Wu BW, Wei B. Liquid state properties and amorphous solidification kinetics of multicomponent $Fe_{50-x}Co_xCr_{14}Mo_{14}C_9B_8Tm_5$ alloys investigated under containerless processing conditions. *Phys Rev E*. 2024;110(3):034612.

doi: 10.1103/physreve.110.034612

14. Mahbooba Z, Thorsson L, Unosson M, *et al*. Additive manufacturing of an iron-based bulk metallic glass larger than the critical casting thickness. *Appl Mater Today*. 2018;11:264-269.

doi: 10.1016/j.apmt.2018.02.011

15. Hooper PA. Melt pool temperature and cooling rates in laser powder bed fusion. *Addit Manuf*. 2018;22:548-559.

doi: 10.1016/j.addma.2018.05.032

16. Yang Z, Markl M, Körner C. Comprehensive numerical investigation of laser powder bed fusion process conditions for bulk metallic glasses. *Addit Manuf*. 2024;81:104026.

doi: 10.1016/j.addma.2024.104026

17. Li B, Yakubov V, Nomoto K, *et al*. Superior mechanical properties of a Zr-based bulk metallic glass via laser powder bed fusion process control. *Acta Mater*. 2024;266:119685.

doi: 10.1016/j.actamat.2024.119685

18. Chen C, Fan Y, Zhang W, *et al*. Tailoring Nano-crystallization in Zr50Ti4Y1Al10Cu25Ni7Co2Fe1 complex multicomponent bulk metallic glass by O doping. *J Non Cryst Solids*. 2021;553:120474.

doi: 10.1016/j.jnoncrysol.2020.120474

19. Chen Y, Zhang D, O'Toole P, *et al*. *In situ* observation and reduction of hot-cracks in laser additive manufacturing. *Commun Mater*. 2024;5(1):84.

doi: 10.1038/s43246-024-00522-3

20. Ansari MA, Crampton A, Garrard R, Cai B, Attallah M. A convolutional neural network (CNN) classification to identify the presence of pores in powder bed fusion images. *Int J Adv Manuf Technol*. 2022;120(7):5133-5150.

doi: 10.1007/s00170-022-08995-7

21. Westphal E, Seitz H. A machine learning method for defect detection and visualization in selective laser sintering based on convolutional neural networks. *Addit Manuf*. 2021;41:101965.

doi: 10.1016/j.addma.2021.101965

22. Mozaffar M, Paul A, Al-Bahrani R, *et al*. Data-driven prediction of the high-dimensional thermal history in directed energy deposition processes via recurrent neural networks. *Manuf Letters*. 2018;18:35-39.

doi: 10.1016/j.mfglet.2018.10.002

23. Chen Z, Mak S, Wu CF. A hierarchical expected improvement method for Bayesian optimization. *J Am Stat Assoc*. 2024;119(546):1619-1632.

doi: 10.1080/01621459.2023.2210803

24. Lei B, Kirk TQ, Bhattacharya A, *et al*. Bayesian optimization with adaptive surrogate models for automated experimental design. *NPJ Comput Mater*. 2021;7(1):194.

doi: 10.1038/s41524-021-00662-x

25. Zhang H, Chen WW, Iyer A, Apley DW, Chen W. Uncertainty-aware mixed-variable machine learning for materials design. *Sci Rep*. 2022;12(1):19760.

doi: 10.1038/s41598-022-23431-2

26. Ament S, Daulton S, Eriksson D, Balandat M, Bakshy E. *Unexpected Improvements to Expected Improvement for Bayesian Optimization*. [arXiv Preprint]; 2025.

doi: 10.48550/arXiv.2310.20708

27. Johnson JE, Jamil IR, Pan L, Lin G, Xu X. Bayesian optimization with Gaussian-process-based active machine learning for improvement of geometric accuracy in projection multi-photon 3D printing. *Light Sci Appl*. 2025;14(1):56.

doi: 10.1038/s41377-024-01707-8

28. Karkaria V, Goeckner A, Zha R, *et al*. Towards a digital twin framework in additive manufacturing: Machine learning and Bayesian optimization for time series process optimization. *J Manuf Syst*. 2024;75:322-332.

doi: 10.1016/j.jmsy.2024.04.023

29. Frazier PI. *A Tutorial on Bayesian Optimization*. [arXiv Preprint]; 2018.

doi: 10.48550/arXiv.1807.02811

30. Laghi V, Palermo M, Bruggi M, Gasparini G, Trombetti T. Blended structural optimization for wire-and-arc additively manufactured beams. *Prog Addit Manuf*. 2023;8(3):381-392.

doi: 10.1007/s40964-022-00335-1

31. Kavas B, Balta EC, Tucker MR, *et al*. *In-situ Controller Autotuning by Bayesian Optimization for Closed-loop Feedback Control of Laser Powder Bed Fusion Process*. Available from: https://arxiv.org/html/2406.19096v1 [Last accessed on 2025 Aug 11].

32. Squires L, Roberts E, Bandyopadhyay A. Radial bimetallic structures via wire arc directed energy deposition-based additive manufacturing. *Nat Commun*. 2023;14(1):3544.

doi: 10.1038/s41467-023-39230-w

33. Yang Z, Markl M, Körner C. Predictive simulation of bulk metallic glass crystallization during laser powder bed fusion. *Addit Manuf*. 2022;59:103121.

doi: 10.1016/j.addma.2022.103121

34. Laws KJ, Miracle DB, Ferry M. A predictive structural model for bulk metallic glasses. *Nat Commun*. 2015;6(1):8123.

doi: 10.1038/ncomms9123

35. Ding J, Ma E. Computational modeling sheds light on structural evolution in metallic glasses and supercooled liquids. *NPJ Comput Mater*. 2017;3(1):9.

    doi: 10.1038/s41524-017-0007-1

36. Xu D, Johnson WL. Crystallization kinetics and glass-forming ability of bulk metallic glasses $Pd_{40}Cu_{30}Ni_{10}P_{20}$ and $Zr_{41.2}Ti_{13.8}Cu_{12.5}Ni_{10}Be_{22.5}$ from classical theory. *Phys Rev B*. 2006;74(2):024207.

    doi: 10.1103/physrevb.74.024207

37. Bordeenithikasem P, Liu J, Kube SA, *et al*. Determination of critical cooling rates in metallic glass forming alloy libraries through laser spike annealing. *Sci Rep.* 2017;7(1):7155.

    doi: 10.1038/s41598-017-07719-2

38. Lázaro-Gredilla M, Titsias MK. Variational Heteroscedastic Gaussian process Regression. In: *Proceedings of the 28th International Conference on International Conference on Machine Learning*. *ICML'11*. Washington, DC: Omnipress; 2011. p. 841-848.

39. Heinonen M, Mannerström H, Rousu J, Kaski S, Lähdesmäki H. *Non-Stationary Gaussian Process Regression with Hamiltonian Monte Carlo*. [arXiv Preprint]; 2015.

    doi: 10.48550/arXiv.1508.04319

40. Jospin LV, Buntine W, Boussaid F, Laga H, Bennamoun M. Hands-on Bayesian Neural Networks -- a Tutorial for Deep Learning Users. *IEEE Comput Intell Mag*. 2022;17(2):29-48.

    doi: 10.1109/mci.2022.3155327

41. Gal Y, Ghahramani Z. *Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning*. [arXiv Preprint]; 2016.

    doi: 10.48550/arXiv.1506.02142

42. Knowles J. ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Trans Evol Comput*. 2006;10(1):50-66.

    doi: 10.1109/tevc.2005.851274

43. Hernández-Lobato JM, Gelbart MA, Adams RP, Hoffman MW, Ghahramani Z. *A General Framework for Constrained Bayesian Optimization using Information-Based Search*. [arXiv Preprint]; 2016.

    doi: 10.48550/arXiv.1511.09422

44. Wong TT. Performance evaluation of classification algorithms by *k*-fold and leave-one-out cross validation. *Pattern Recognit*. 2015;48(9):2839-2846.

    doi: 10.1016/j.patcog.2015.03.009

45. Vehtari A, Gelman A, Gabry J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Stat Comput*. 2017;27(5):1413-1432.

    doi: 10.1007/s11222-016-9696-4