Journal of Laser Applications

RESEARCH ARTICLE | SEPTEMBER 23 2025

Bayesian optimization for extreme high-speed laser material deposition

Special Collection: 9th International Congress on Laser Advanced Materials Processing

Max Gero Zimmermann ✉ ; Johanna Menn ; Christopher Ullmann ; Viktor Glushych ; Thomas Schopphoven ; Wilhelm Meiners ; Sebastian Trimpe ; Carlo Holly

J. Laser Appl. 37, 042013 (2025)
https://doi.org/10.2351/7.0001906

Journal of Laser Applications ARTICLE pubs.aip.org/lia/jla

Bayesian optimization for extreme high-speed laser material deposition

Cite as: J. Laser Appl. 37, 042013 (2025); doi: 10.2351/7.0001906
Submitted: 30 June 2025 · Accepted: 1 September 2025 ·
Published Online: 23 September 2025

Max Gero Zimmermann,[1,a)] Johanna Menn,[2] Christopher Ullmann,[1] Viktor Glushych,[1] Thomas Schopphoven,[1] Wilhelm Meiners,[1] Sebastian Trimpe,[2] and Carlo Holly[1,3]

AFFILIATIONS

[1]Fraunhofer Institute for Laser Technology ILT, Aachen 52070, Germany
[2]Institute for Data Science in Mechanical Engineering, RWTH Aachen University, Aachen 52062, Germany
[3]Chair for Technology of Optical Systems TOS, RWTH Aachen University TOS, Aachen 52074, Germany

Note: This paper is part of the Special Topic on the 9th International Congress on Laser Advanced Materials Processing.
[a)]**Author to whom correspondence should be addressed; electronic mail:** max.zimmermann@ilt.fraunhofer.de

ABSTRACT

Extreme high-speed laser material deposition, also known by the acronym EHLA, enables metallic coatings of different thicknesses at deposition speeds of up to several hundred meters per minute and deposition rates of several kilograms per hour. Against other deposition welding processes, EHLA offers significant advantages in terms of lower heat input and higher precision, making it a valuable option for processing materials, which are considered hard-to-weld. Despite its advantages, the highly nonlinear interdependencies of multiple influencing variables require precise control and tuning of the parameters and challenge the process development, making it time and cost expensive. In the absence of an accurate process model for large parameter spaces, model-based optimization is currently not feasible, such that current development requires extensive experimentation and expert knowledge. To overcome these challenges, an adaptive process development approach for the key process parameters, such as laser power and powder flow rate, based on Bayesian optimization (BO) is proposed. BO employs probabilistic models trained on experimental data to systematically explore the parameter space and predict the optimal settings in terms of a target variable. The investigations show that the sample-efficient, data-driven method effectively accelerates the development of suitable process parameters and drastically reduces the need for extensive empirical testing and expert knowledge.

Key words: extreme high-speed laser material deposition (EHLA), laser coating, parameter development, surface analysis, Bayesian optimization

I. INTRODUCTION

Extreme high-speed laser material deposition (EHLA) is a resource- and cost-efficient laser-cladding technique used to coat cylindrical metallic components, yielding higher powder utilization than conventional laser material deposition (LMD) methods (Sec. I A). However, industrial deployment remains challenging: EHLA requires tuning over 19 interacting parameters, and manual, iterative optimization by experts, which is both time consuming and expensive.[1,2] Automatic parameter tuning can, therefore, significantly reduce development costs and, as a result, increase the feasibility of EHLA for industrial applications. One key objective is to increase the coating thickness while respecting process productivity and surface-quality constraints. As simulations are not always available, process parameter tuning must be done in the real process. Unsuccessful parameter sets may result in the detachment of the coating layer ("crashed" experiments), and each experiment, along with thickness and quality measurements, takes considerable time. To reduce setup times, process optimization should only require a few experiments. In summary, the automatic optimization of EHLA faces the following key challenges:

C1. On-machine experiments.
C2. Quality constraints.
C3. Crashed experiments.
C4. Optimization with batches of experiments.

Existing EHLA tuning is mostly manual, and currently, no automatic method addresses all four challenges simultaneously. Bayesian optimization (BO) (Sec. I B) is a black-box technique well-suited for expensive experiments with constraints, crashed experiments, and batch optimization. It has been successfully applied in other laser-based processes (Sec. I C) and enables data-efficient optimization on real processes without the need for a simulation.

In this work, we first define an industrially relevant problem for parameter optimization in EHLA (Sec. I D). We propose a constrained, crash-aware, and batch BO framework for EHLA and demonstrate its application to tune two key parameters of the process (Sec. II). We construct a surrogate model to evaluate algorithmic variants and identify a BO configuration (Secs. III A–III D). Finally, we validate it on a real EHLA process, achieving semi-automatic process optimization for the first time (Sec. III E).

A. EHLA

The EHLA process represents an advancement of conventional laser cladding, also referred to as laser-based directed energy deposition or LMD. It is an additive manufacturing technique in which a material is selectively deposited onto a substrate and simultaneously melted by a laser. EHLA fundamentally differs from classical LMD processes in the way the laser is coupled to create depositions.[3] In EHLA, the metal powder particles are fully melted above the melt pool.[1] Thus, unlike conventional LMD, where time is required to generate and sustain a melt pool on the component surface in which the powder must be completely melted, EHLA eliminates that delay.[3] The high proportion of already molten powder particles forms the melt pool directly, allowing the surface speed (i.e., the relative velocity between the cladding optics and the part) to be increased to up to 500 m/min (standard LMD speeds are usually up to 2 m/min).[4] Consequently, EHLA can produce thin coatings in the range of 25–250 μm.[4,5] Furthermore, the high surface speed leads to a reduced heat-affected zone, minimizing thermal distortion and making the process particularly suitable for thin-walled or geometrically sensitive components.[1] In addition, the complete melting of the powder particles above the substrate surface ensures that, upon impact with the workpiece, the powder particles flow directly into the melt pool, achieving a powder utilization rate of >95% (classical LMD achieves >80%).[2] Combined with the high deposition rates, EHLA offers a cost-effective and resource-efficient alternative to LMD for functional surface protection.[1,6]

B. Bayesian optimization

BO is a black-box optimization method that effectively finds the optimum of an unknown (objective-) function by iteratively querying the function.[7] In BO, the objective function is modeled with a probabilistic surrogate model, often a Gaussian process (GP). This model provides a current prediction of the function as well as uncertainty estimates. To find the optimum of the objective function, the uncertainty estimates are leveraged by the acquisition function, which suggests the next point to query. In the process parameter tuning setting, a function evaluation corresponds to performing one experiment with a certain parameter configuration.

C. Related work

The optimization of parameters in EHLA is complex due to nonlinear interactions, which have led to a wide variety of approaches. These range from simple one-factor experiments to statistical methods such as response-surface methodology, and machine learning methods like artificial neural networks or support vector machines. A comprehensive overview is provided by Wang *et al.*[8] Another approach to process parameter optimization is the use of simulation methods. Numerical simulations, for example, those based on the finite element method, enable the prediction of temperature fields, melt-pool geometries, stress distributions, and microstructure evolution.[9] Nevertheless, experimental reference points must be generated for the simulations, which serve as input data or validation points.[10] Simulation enables a deeper understanding of the process, allowing for knowledge-based process optimization but not automated process optimization.[11]

In contrast to these works, our BO approach aims to optimize the process parameters directly on the machine. This avoids the sim-to-real gap and enables fast process optimization of process configurations and materials for which simulations are unavailable.

BO has already been applied in various laser-based processes for process optimization.[12,13] Kavas *et al.* use BO to adaptively optimize the parameters of an in-layer proportional integral (PI) controller in the laser powder bed fusion (LPBF) process, thereby avoiding overheating in critical geometries and increasing process stability.[14] To optimize multiple objectives in LPBF simultaneously (hardness and porosity), Chepiga *et al.* employ a diversity-guided efficient multi-objective optimization (DGEMO) algorithm, achieving a pareto-optimal configuration within six iterations.[15] Karkaria *et al.* utilize BO as part of a digital-twin framework to determine the optimal laser power for the LMD process using a parameterized time-series generator and a GP surrogate model, thus minimizing simulation effort and heat input while attaining desired mechanical properties.[16] Sousa *et al.* demonstrate a "human-in-the-loop" approach for LMD, showing how real-time *in situ* monitoring combined with multi-objective Bayesian optimization can jointly optimize melt-pool dimensions and stability. Their results indicate that, with a response-surface-driven BO workflow, target metrics can be efficiently achieved and the number of required single-wire trials can be significantly reduced.[17]

Our paper is the first application of BO in EHLA tackling its unique challenges, such as constraints, experimental crashes and batch-setting, which have not been evaluated in this context before.

D. Problem definition

We aim to find the optimal parameters x^* to maximize coating thickness $t: S \to \mathbb{R}$ within an n-dimensional parameter space $S \subset \mathbb{R}^n$,

$$x^* \in \arg\max_{x \in S} t(x). \tag{1}$$

Here, $t(x)$ denotes the function describing coating thickness as a function of the process parameters $x \epsilon S$. Simultaneously, when maximizing coating thickness, quality criteria, such as crack-free or pore-free surfaces, are imposed on the coating, resulting in

parameter constraints of the form $c_i(x) \leq y_i$ or $c_i(x) \geq y_i$, where y_i represents the threshold value for each constraint. The parameter space S can be further restricted by physical infeasibilities. For example, sufficient energy must be supplied to fully melt the fed powder material. If the energy is too low, no coating can be generated. Consequently, the parameter space S is limited to a feasible parameter space S_s, with $S_s \subseteq S$. The complementary set of parameters that yield impossible coatings is denoted $S_c = S \backslash S_s$. The assignment of parameters to the corresponding subsets can be described by

$$a(x) = \begin{cases} 1, & x \in S_S \\ 0, & x \in S_C \end{cases}. \tag{2}$$

Thus, a general parameter-optimization problem for maximizing a requirement while complying with constraints can be formulated as

$$\begin{aligned} &\max_{x \in S} t(x), \\ &s.t.\ c_i(x) \leq y_i,\ a(x) = 1. \end{aligned} \tag{3}$$

II. METHODOLOGY AND MATERIALS

As an approach to solve the optimization problem [Eq. (3)], we propose batch BO with constraints and crash constraints. Our goal is to develop an optimization approach to be used directly on the machine (key challenge C1). Therefore, we need to define our experimental setup, a measurement procedure to quantify objectives and constraints, and a suitable algorithmic configuration.

As this is the first application of BO for EHLA, we identify two parameters, powder mass flow and laser power, that directly influence the coating thickness without reducing economic efficiency, as intuitively higher powder mass flow leads to higher coating thickness and melting more powder a higher laser power is needed. We choose parameter bounds for the laser power $P_L \in [3, 8\ \text{kW}]$ and for the powder mass flow $\dot{m}_P \in \left[15, 120\ \frac{g}{\text{min}}\right]$, resulting in the two-dimensional search space $S \subset \mathbb{R}^2$. We evaluate our algorithm in this setting, while keeping the other parameters constant as described in Table I.

To be able to optimize the coating thickness, we need to conduct an experiment and measure it. For this, a VR-5000 3D profilometer from Keyence is used. With this instrument, coating thickness can be measured over a large area ($70.000 \times 25.000\ \mu\text{m}^2$ per coating), and the accompanying Keyence software can be used to determine various thickness metrics, such as the minimum and maximum thickness. An exemplary evaluation is shown in Fig. 1. We use the minimal measured thickness t_{min} since this represents the technically relevant thickness after post-processing.[18]

As a surrogate model for the unknown objective function, a GP was used within BO. As an acquisition function, expected improvement (EI) is used, which maximizes the expected increase over the current best result t_{best} such that

$$EI(x) = \mathbb{E}[t(x) - t_{\text{best}}]. \tag{4}$$

While maximizing the coating thickness by the two parameters in the search space is intuitively possible, as high powder mass flow and high laser power lead to high coating thickness. However, not all coatings are usable as they do not fulfill coating quality constraints (C2). In order to be able to make a statement about the coating quality, publications show that surface analysis correlates with coating quality.[19] By using surface analysis instead of metallography, there is a time saving of approximately 20 min per coating. To reduce the experiment time, we also use surface analysis to analyze coating quality with a Zygo Nexview NX2 white-light interferometer (WIM). With the WIM, the central area of each coating is imaged at 14×, 55×, and 200× magnification. The inflow and outflow areas of the coating are not examined in detail. To determine the surface roughness Ra from the WIM data, the Fraunhofer ILT's internal software RCALCLL is used.[20] This software applies bandpass filters to visualize different roughness scales. A visualization of various roughness scales is provided in Fig. 2. The surface analysis yields 20 different metrics that can be used as constraints. Therefore, we determine in Sec. III A whose metric combination is suitable to classify ok and nok coatings and investigate the influence on optimization.

TABLE I. EHLA parameters used.

Parameter	Value
Laser power	3–8 kW
Powder feed rate	15–120 g/min
Process speed	70 m/min
Feed rate	300 μm/rev
Feed gas rate	12 l/min
Shielding gas rate	12 l/min
Distance from nozzle tip to substrate	15.000 μm
Coating width	20.000 μm
Distribution ring to core	4%
Focal length collimation	30.000 μm

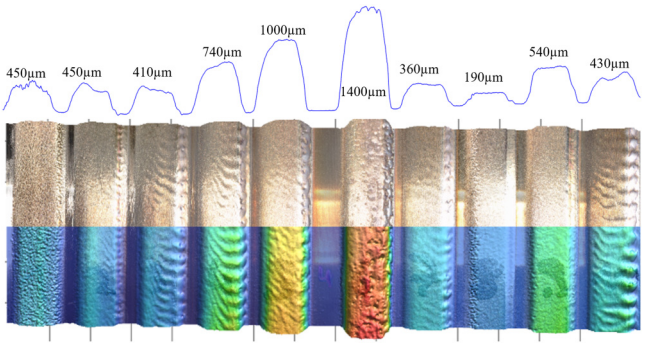

FIG. 1. Minimum coating thickness (top), surface images (middle), and waviness and thickness variation maps (below) for EHLA coatings produced with different parameter settings, measured over $70.000 \times 25.000\ \mu\text{m}^2$ using a Keyence VR-5000 3D profilometer. Colors represent thickness in μm; cooler colors indicate thinner regions and warmer colors indicate thicker regions.

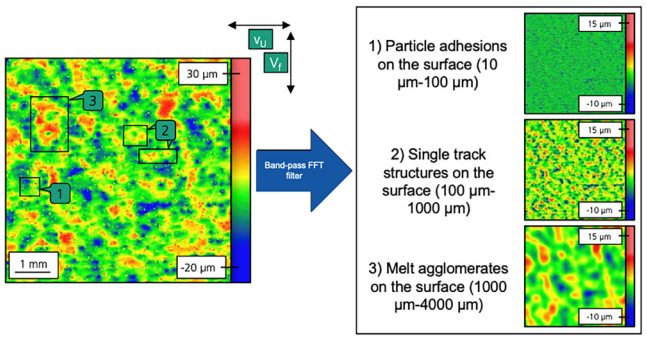

FIG. 2. Evaluation of WIM data using RCALCLL software. The roughness is divided into different areas using bandpass filters.

To include the constraints in the BO framework, we use an adjusted version of EI, constrained EI (CEI).[21] In this setting, we model the objective function t, as well as the constraint functions all as separate GPs. The GPs of the constraint functions provide an estimate of how likely it is, that a point is feasible or not. This probability of feasibility P_t is multiplied by the standard acquisition function of EI, to suggest the next query point, resulting in

$$\mathrm{CEI} = \mathbb{E}\left[t_n^* - \min_{i=1,\dots,q} t(x_i)^+\right] P_t(x). \tag{5}$$

If a parameter set leads to process failure, for example, if the coating does not adhere, no surface analysis can be performed (C3), and thus, no data can be returned to the BO algorithm. Therefore, we implement crash constraints in addition to the standard constraints.[22] Nonevaluatable parameter sets are returned to the BO algorithm as virtual observations to steer the algorithm away from unstable or infeasible regions of the parameter space. As BO can react sensitively to the data provided, the virtual observations must be selected so that they are not unrealistically far from the real measured data but remain unfavorable enough to keep the search away from crash regions. Following Stenger and Abel,[23] we define virtual observations at previously observed crash locations $\hat{x}$ as

$$\hat{y}_i = \max\left(\mu(\hat{x}),\ \mu(x_*)\right) + \beta\sqrt{k_D(\hat{x})}, \tag{6}$$

where $\mu(\cdot)$ denotes the surrogate model's posterior mean, x_* is the current parameterization, $k_D(\cdot)$ is the posterior variance (under dataset D) so that $\sqrt{k_D(\cdot)}$ is the posterior standard deviation, and $\beta > 0$ is a penalty parameter. Rohr *et al.* show that $\beta = 3$ is an empirically effective choice and is, therefore, also used in this paper.[22] Through the investigations in Sec. III D, it can be evaluated how crash constraints affect the approximation behavior of the BO algorithm.

In EHLA, the experiment itself is fast (30 s). However, the measurement procedure is very time consuming as process overheads arise during each analysis, parts have to be unclamped, set up on measurement stations, and reintegrated into the test setup after analysis. If multiple coatings are analyzed simultaneously, this measurement overhead can be reduced (C4). To deal with this special requirement, we employ batch optimization in BO as it offers the advantage of evaluating multiple data points in parallel per iteration. This can increase optimization efficiency.[24] On the other hand, each coating experiment incurs costs that should be minimized. By investigating batch size (see Sec. III B), a trade-off can be established between overhead, optimization efficiency, and experimental costs. For batch BO, the acquisition function changes to the multipoint version q-EI. In accordance with Ament *et al.*,[25] we use a logarithmic transformation on the output, as this yields improved optimization. In combination with constraints, we get the acquisition function

$$q - \mathrm{Log\ CEI}(x) = \log\left(\mathbb{E}\left[t_n^* - \min_{i=1,\dots,q} t(x_i)^+\right]\right) + \log P_t(x), \tag{7}$$

where P_t is the probability of the constraints c_i being fulfilled. Combined with the other extensions, this leads to a BO framework for EHLA with constraints, crash constraints, and batch optimization.

Currently, there are no state-of-the-art studies on constraints with corresponding threshold values, the effectiveness of crash constraints, batch size, or the number of initial data points for process development and optimization in EHLA using BO. Therefore, we examine the influence of these factors in this work. The investigations are carried out using a surrogate model, hereafter referred to as GP_{EHLA}. GP_{EHLA} is a GP based on 39 pre-experiments with associated analyses. The findings were integrated into the BO framework. The BO algorithm is then validated for process development and optimization in the EHLA using new experiments.

The experimental investigations were performed on an EHLA system by Hornet, which is equipped with a TruDisk 12001 laser, a 200/700 μm fiber, and BrightLine Weld technology from Trumpf. The laser beam is focused onto the substrate surface using a Trumpf BEO-D70 optics with motorized collimation, while the metal powder is fed concentrically through an HD-Sonderoptiken High-No 5.0 powder nozzle. An Oerlikon Twin 150-LC unit serves as the powder delivery system. As an application example, an EHLA coating for corrosion prevention was chosen. An unalloyed S355 steel tube was used as the substrate and coated with corrosion-resistant 316L stainless steel powder in the particle-size range of 20–65 μm. Each coating has a width of 20.000 μm. The BO algorithm was implemented in Python using the BoTorch library from Meta.

III. RESULTS

The results of this work are presented in accordance with the outlined procedure. First, the investigations on the surrogate model GP_{EHLA} are presented. In this context, the constraints and corresponding threshold values are identified (see Sec. III A). Next, the batch size at which BO approximates high coating thicknesses most rapidly is investigated (see Sec. III B). Furthermore, the influence of the number of initial data points on GP_{EHLA} was examined (see Sec. III C). Next, the effect of crash constraints within the framework was validated (see Sec. III D). The insights gained were then implemented in the BO algorithm, and process development

in EHLA using the developed BO framework was validated through new experiments (see Sec. III E).

A. Constraints

In the first step, the constraints and their threshold values for distinguishing between ok and nok coatings were investigated. To identify the constraints and their threshold between ok and nok, the 39 pre-experiments are first evaluated by five different experts and classified into 12 ok and 27 nok coatings, with all experts arriving at the same result, demonstrating a very high reliability of the classification. Subsequently, the 39 coatings were measured, yielding 20 distinct metrics from the 3D profilometer and WIM data (see Sec. II). Each metric was mapped to the corresponding expert decision. This analysis revealed that, for every single metric, the minimum and maximum values for ok and nok classifications overlapped. Therefore, no single metric alone could reliably separate the two classes, and a combination of two metrics has to be used. For each metric, we defined a preliminary threshold as the maximum value observed among all ok-classified coatings plus an additional 10% margin. This approach ensures that, first, all coatings classified ok by the experts are included; second, a safety margin is provided to account for process variability; and third, thresholds remain aligned with the resolution of the applied measurement instruments. Next, all possible combinations of the 20 metrics with the individual threshold (190 possibilities) were formed and re-evaluated against the expert decisions. From these, six combinations and their threshold were identified that yielded a clear separation between ok and nok coatings. To determine which constraint combination enabled the BO algorithm to most efficiently maximize coating thickness, five BO configurations (each using a different set of constraints) were tested over ten iterations, by 30 different seed data. The average layer thickness calculated from the 30 seed data for the different BO configurations is shown in Fig. 3 for each iteration.

It should be noted that all six BO configurations exhibit comparable trends in coating thickness development over iterations, indicating that any of the constraint combinations would be suitable. In the first three iterations, however, the BO configuration employing the constraints *<1280 µm + coating thickness difference* approximated the largest coating thicknesses most rapidly, but in subsequent iterations, it performed worse than the other BO configurations. After ten iterations, the BO configuration with the constraint combination *particle + coating thickness difference* achieved the highest coating thickness of 320 μm. In this work, the constraints surface roughness *particle + coating thickness difference* were, therefore, selected as c_1 and c_2, respectively. The roughness metric *particle* refers to the roughness spectrum in the range of 10–80 μm, capturing particle adhesions on the order of the powder particle size (20–65 μm). The constraint *coating thickness difference* accounts for both the variation in thickness across the coating surface and the waviness within a single layer. An example of this can be seen in the fourth coating from the left in Fig. 1. The threshold values for the constraints were determined as described above by aligning them with the experts ok and nok classifications, adding a safety margin of 10% of the respective maximum value. The resulting thresholds are $c_1 = 2, 6\ \mu$m and $c_2 = 100\ \mu$m.

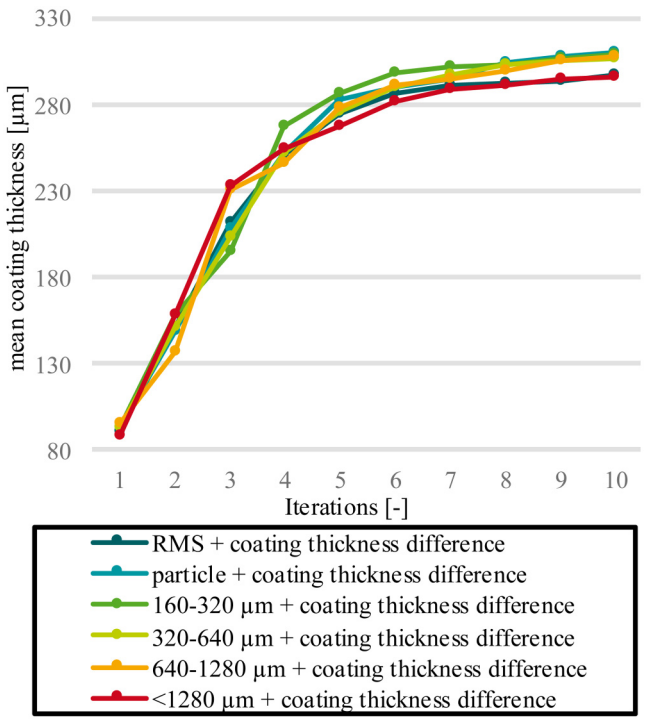

FIG. 3. Mean layer thickness achieved over ten BO iterations for different constraint combinations. Each line represents a BO configuration using one constraint pair (see legend), averaged over 30 trials.

B. Batch size

To address the trade-off between experimental and analysis costs, BO is implemented in this work as a batch optimization. Due to the limited measurement capacity of the 3D profilometer and WIM, up to six coatings can be measured per analysis, so the maximum batch size is six. The minimum batch size is one experiment. Since the influence of the number of initial data points on the batch size is unknown, the study is conducted with varying numbers of initial data points. Corresponding to the maximum batch size, the maximum number of initial data points is also set to six, while the minimum number of initial data points is two. For each batch size, five BO configurations with different numbers of initial data points are run for ten iterations each (totally, 30 BO experiments). Afterward, for each batch size, the average coating thickness achieved across the different initial data-point sets is calculated. The results are shown in Fig. 4.

All BO configurations start with a comparable average coating thickness of approximately 100 μm. In iterations one through three, there is a clear ranking of achieved coating thickness by the batch size (see Fig. 4, iteration 3: batch size 1 at 150 μm, increasing to batch size 6 at 230 μm). This ranking corresponds to the number of data points in the BO configurations (iteration 3: batch size $1 \hat{=} 5 - 9$ data points; batch size $6 \hat{=} 20 - 24$ data points). Starting from the fourth iteration, this ranking changes and the BO configurations

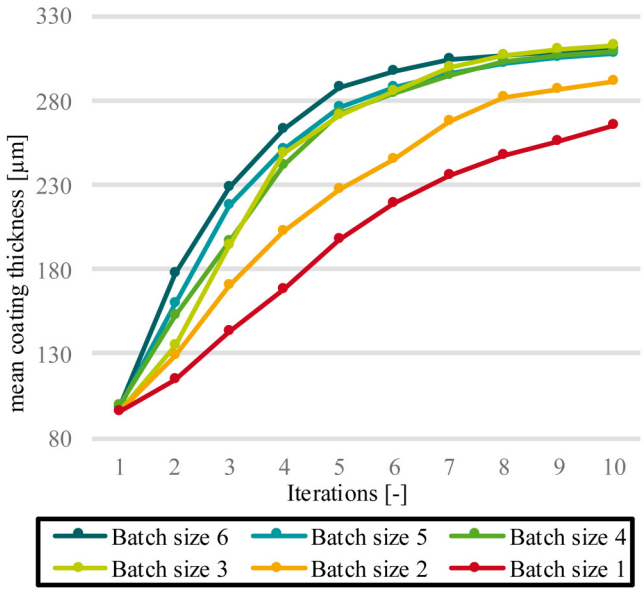

FIG. 4. Effect of the batch size on BO performance, shown as the mean of the maximum coating thickness obtained at each iteration for different numbers of initial data points.

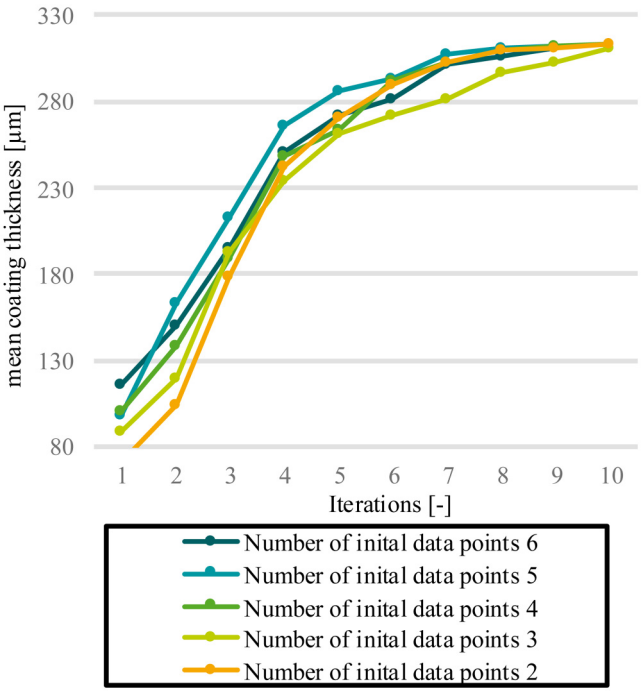

FIG. 5. Influence of the number of initial data points on BO performance with a fixed batch size of three. All configurations converge to similar final coating thickness (~315 μm), but starting with five initial data points yields slightly higher thicknesses from iteration 2 onward.

with batch size 3 achieve a higher coating thickness than the model with batch size 4. After six iterations, the difference between batch size 6 and batch size 3 is 12 μm, which decreases in the seventh iteration to 4 μm. For the real experiment, this means that with batch size 6, we would have conducted 18 more experiments by the end of six iterations compared to batch size 3. After ten iterations, the BO configuration with batch size 3 reaches the highest coating thickness of 313 μm (with batch size 6 reaches 310 μm). Therefore, a batch size 3 offers an acceptable trade-off between experimental cost, analysis cost, and BO accuracy and was used for the subsequent procedure.

C. Initial data points

To investigate the BO algorithm in terms of the number of initial data points, an additional experiment was conducted using the surrogate model GP_{EHLA}. As described for batch size, the number of initial data points was chosen between two and six. This results in five different BO configurations. For each configuration, ten iterations with six different starting values were run with a batch size of three. Figure 5 shows the results of the average layer thickness achieved per iteration.

After ten iterations, all configurations reach comparable coating thicknesses of approximately 315 μm. Figure 5 indicates that from the second through the tenth iteration, the BO configurations with five initial data points achieve the highest coating thickness, comparable to the BO configuration with other numbers of initial data points. This configuration was implemented in the framework for further processing.

D. Crash constraints

In addition to constraints, the performance of the BO algorithm is examined when proposed parameters lead to nonadherent and, thus, nonevaluable coatings (referred to as $\mathrm{BO}_{\mathrm{withoutCC}}$). In contrast, the BO algorithm with crash constraints (hereafter $\mathrm{BO}_{\mathrm{withCC}}$; see Sec. II) was also evaluated. For each algorithm type ($\mathrm{BO}_{\mathrm{withoutCC}}$ and $\mathrm{BO}_{\mathrm{withCC}}$), 39 experiments were conducted using the same initial parameter sets. The $\mathrm{BO}_{\mathrm{withoutCC}}$ model received no feedback for nonevaluable parameter suggestions, resulting in fewer data points available to fit the GP surrogate. The results are presented as box plots in Fig. 6. From the second through the sixth iteration, $\mathrm{BO}_{\mathrm{withCC}}$ achieves, on average, higher coating thicknesses than $\mathrm{BO}_{\mathrm{withoutCC}}$. Furthermore, up to the fifth iteration, the interquartile range for $\mathrm{BO}_{\mathrm{withCC}}$ is smaller than for $\mathrm{BO}_{\mathrm{withoutCC}}$. Beyond the fifth iteration, however, $\mathrm{BO}_{\mathrm{withoutCC}}$ exhibits a smaller interquartile range than $\mathrm{BO}_{\mathrm{withCC}}$. This behavior can be attributed to $\mathrm{BO}_{\mathrm{withoutCC}}$ acquisition function focusing the search on a narrower region of the parameter space. Batch optimization further amplifies this effect, since each batch contained at least one evaluable point; in a single-batch scenario, $\mathrm{BO}_{\mathrm{withoutCC}}$ would perform worse.

E. Experimental validation

The preliminary studies on the surrogate model GP_{EHLA} (Secs. III A–III D) have determined which constraints and

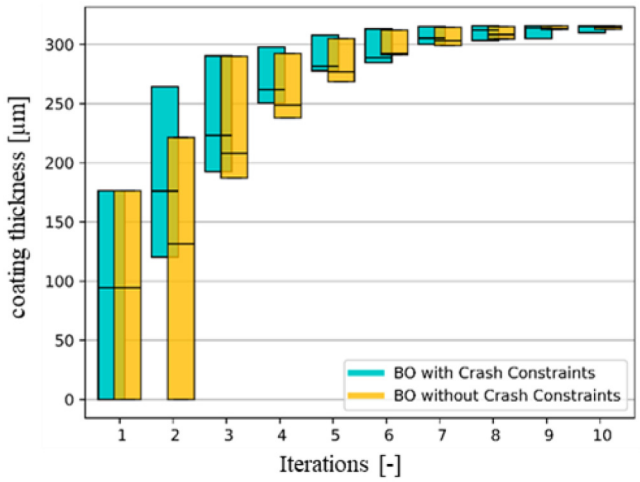

FIG. 6. Distribution of coating thicknesses achieved by BO with crash constraints (BO_{withCC}, turquoise color) and without crash constraints ($BO_{withoutCC}$, yellow color) across 39 experiments. Boxes show interquartile ranges. BO_{withCC} achieves higher median thicknesses in early iterations and maintains more consistent performance until iteration 5.

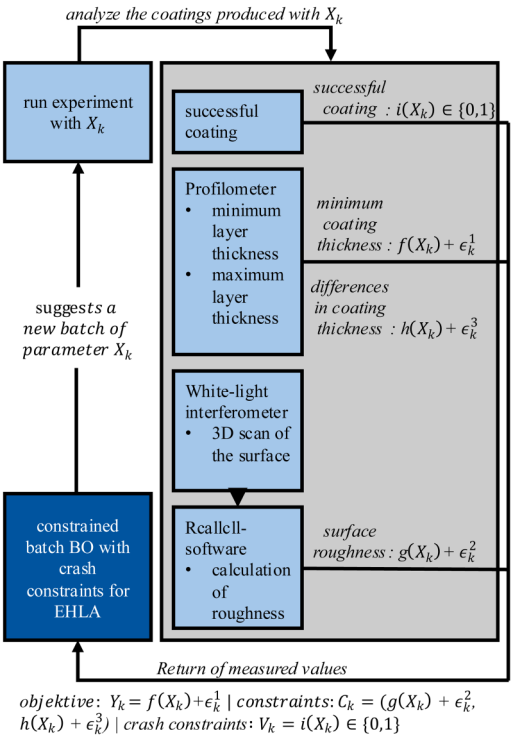

FIG. 7. Schematic of the experimental validation workflow for the BO framework. The process includes initialization with random parameter sets, coating deposition, measurement of objective and constraints, BO model update, and proposal of new parameter sets.

threshold values, which batch size, and how many initial data points enable the BO algorithm to maximize coating thickness in as few iterations as possible. Based on these findings, the framework depicted in Fig. 7 is used to validate BO-driven parameter development in EHLA. A stop criterion has been integrated into the framework so that the loop in Fig. 7 is to be run through one last time as soon as a layer thickness of >500 μm is reached while complying with the constraints. The last loop was integrated to check whether the layer thickness could be increased further. For validation with real experiments, the BO algorithm is newly initialized. This means that the BO validation model has not yet received any data points from preliminary studies and that the GP surrogate model within BO has not yet been constructed. As described in Sec. III C, five new randomized initial parameter sets are generated. EHLA coatings are then performed and measured, and the resulting data are fed into the BO algorithm. Subsequently, the framework shown in Fig. 7 was executed for four iterations with a batch size of three.

In total, 17 EHLA coatings were performed and analyzed, of which three exhibited delamination failures (see Fig. 8, white outlines) and were, thus, classified as crash constraints. In 12 coatings, at least one constraint threshold was exceeded, resulting in five coatings being classified as ok. The maximum ok coating thickness achieved was 557 μm, while the maximum nok thickness was 775 μm. The required layer thickness of >500 μm was achieved in the third iteration.

Figure 8 shows the process map of the GP model after the seed experiments and four conducted iterations. The orange dots represent the BO model's prediction for a fifth iteration, which was not performed. In the first panel, the achieved coating thickness is plotted, in the middle panel, the constraint c_2 (coating thickness difference), and in the last panel, the constraint c_1 (surface roughness). The initial parameter sets and those proposed by the BO algorithm are color-coded for each iteration.

It should be noted that the parameter suggestions in iterations 1 and 2 are still widely dispersed across the parameter space. From the third iteration onward, parameter sets are increasingly proposed within a laser power range of 7000–8000 W and a powder mass flow range of 80–100 g/min. For closer analysis, Fig. 9 presents a zoom into this region. The color scale has been adjusted relative to Fig. 8 to highlight the constraint thresholds more clearly. The constraints are satisfied for $c_1 < 2.6\,\mu$m and $c_2 < 100\,\mu$m (see the green-yellow area in Fig. 9). Within this zoomed view, it can be seen that with each successive iteration, proposed parameter sets shift toward regions of higher expected coating thickness (yellow area in the top panel of Fig. 9). Moreover, in iteration 2 (red points in Fig. 9), either c_1 or c_2 is violated. In iterations 3 and 4, one parameter set per iteration triggers the crash constraint (see the white-outlined blue and black points in Fig. 8), while two parameter sets per iteration satisfy both constraints (blue and black points in Fig. 9). Between iterations 3 and 4, the parameter sets are adjusted such that higher coating thicknesses become attainable. For the prediction of iteration 5, the BO algorithm attempts to further increase the coating thickness. However, every proposed parameter set violates at least one constraint threshold.

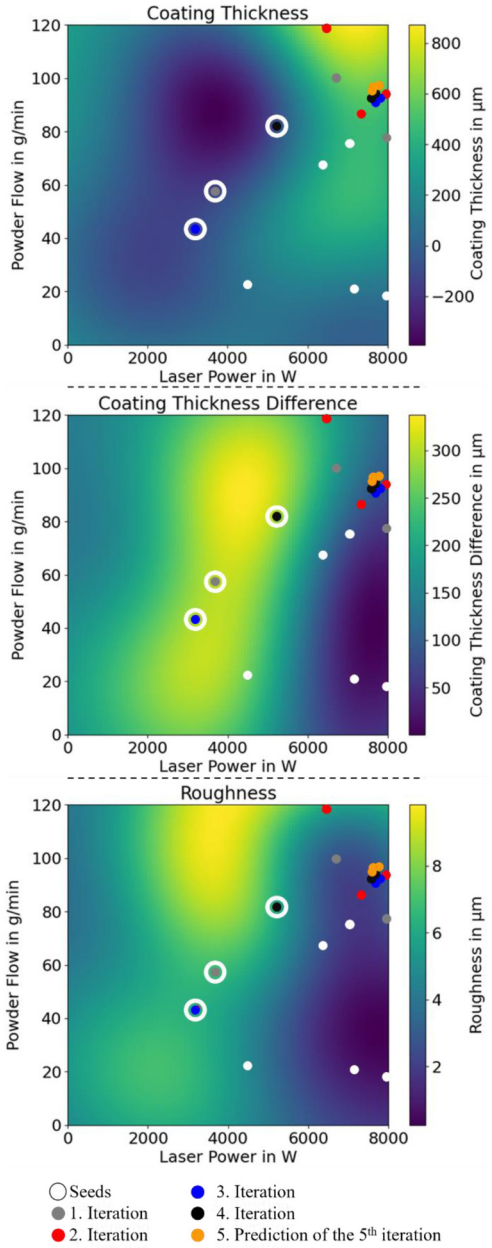

FIG. 8. GP surrogate model predictions after five BO iterations (seed experiments + four iterations) in the validation experiment. Top: predicted coating thickness; middle: predicted coating thickness difference; below: predicted roughness. Colored dots indicate parameter sets tested in each iteration. White outlines mark crash cases (nonadherent coatings).

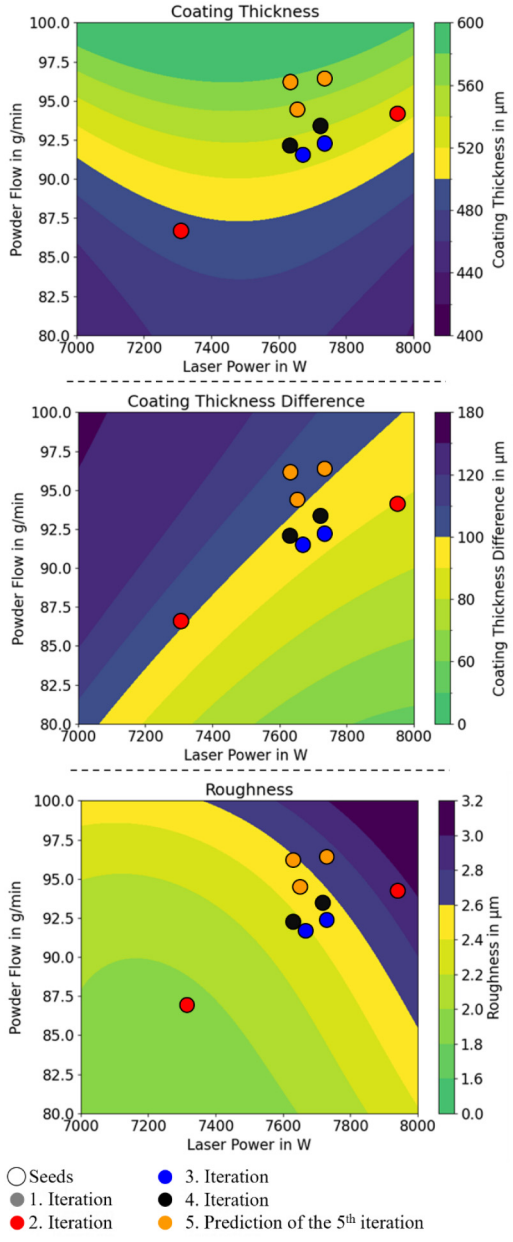

FIG. 9. Zoomed view of the high-thickness, constraint-satisfying region from Fig. 8. Color scales highlight the constraint thresholds (coating thickness difference = 100 μm, roughness = 2.6 μm). Points from later iterations (black) cluster near the optimal region (7712 W laser power, 95 g/min powder flow), achieving up to 557 μm thickness while satisfying both constraints.

By overlaying the process maps of the constraints from iteration 4, it can be observed that the maximum expected coating thickness of 560 μm is achieved under compliance with the constraints at $P_L = 7712$ W and $\dot{m}_P = 95$ g/min. In iterations 3–5, the BO algorithm clearly navigates along the constraint thresholds in its attempts to maximize coating thickness. Notably, in iteration 4, a parameter set of $P_L = 7722$ W and $\dot{m}_P = 94,43$ g/min yielded a coating thickness of 557 μm, corresponding to a theoretical powder efficiency of 100%.

IV. CONCLUSION

In this work, we developed a BO framework and applied BO for the first time to maximize coating thickness in the EHLA process under constant economic conditions and quality requirements. Initially, a GP surrogate model (GP_{EHLA}) was constructed based on 39 pre-experiments, which enabled the identification of critical constraints and their thresholds for ensuring surface quality. Subsequent experiments determined the batch size, and the number of initial data points required for BO to achieve maximal coating thickness in the parameter space with minimal experimental effort. The framework presented in Fig. 7, thus, facilitates largely automated parameter optimization without the need for extensive expert knowledge.

The model was then validated in 17 real EHLA coatings, resulting in the identification of a parameter set yielding a coating thickness of 557 μm while satisfying all quality constraints, more than twice the conventional EHLA thickness of 250 μm. Furthermore, compared to traditional process developments in the 39 pre-experiments, it was demonstrated that BO achieves higher layer thicknesses with fewer experiments. The BO framework is run four times until the stopping criterion is reached. The predictions for a fifth iteration indicate a potentially higher layer thickness. Overall, this study demonstrates that BO is a powerful tool for supporting and accelerating EHLA process development.

V. OUTLOOK

The results demonstrate the successful parameter development for maximizing coating thickness in the EHLA process. This work suggests several avenues for future research to further streamline process development. One approach should focus on extracting relevant quality features directly from EHLA process data, paving the way for a closed-loop BO implementation. The sensitivity of the threshold values of the constraints should also be considered further with regard to the performance of the BO model. The stop criteria must also be considered here. In this work, it can be seen that although the required layer thickness has already been achieved, there are parameters where even higher layer thicknesses exist. Moreover, BO can be extended to handle more than two parameters simultaneously when optimizing a target metric. Although this increases the search space, it also enables the identification of additional nonlinear interactions among parameters. An extension could involve not only integrating more parameters but also multiple optimization objectives, such as hardness, crack count, or porosity, resulting in a multi-objective BO. It remains to be investigated whether a constraint-based optimization or a multi-objective BO yields greater efficiency in parameter development. While BO is a powerful tool for process optimization, questions remain regarding the impact of disturbances, such as variations in powder batches or machine wear, on the BO algorithm. Furthermore, research is needed to determine how learned process–parameter relationships can be transferred to other material systems in order to further reduce experimental costs.

ACKNOWLEDGMENTS

The authors thank David Stenger for helpful insights and discussions. This work was funded, in part, by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC-2023 Internet of Production—No. 390621612 and within the Helmholtz School for Data Science in Life, Earth and Energy (HDS-LEE).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Max Zimmermann: Conceptualization (equal); Formal analysis (equal); Project administration (equal); Supervision (equal); Validation (equal); Writing – original draft (equal). **Johanna Menn:** Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Writing – review & editing (equal). **Christopher Ullmann:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Methodology (equal); Software (equal); Visualization (equal); Writing – review & editing (equal). **Viktor Glushych:** Project administration (equal); Supervision (equal); Writing – review & editing (equal). **Thomas Schopphoven:** Project administration (equal); Writing – review & editing (equal). **Wilhelm Meiners:** Supervision (equal); Writing – review & editing (equal). **Sebastian Trimpe:** Funding acquisition (equal); Supervision (equal); Writing – review & editing (equal). **Carlo Holly:** Funding acquisition (equal); Supervision (equal); Writing – review & editing (equal).

REFERENCES

[1]T. Schopphoven, A. Gasser, and G. Backes, "EHLA: Extreme high-speed laser material deposition," Laser Tech. J. **14**, 26–29 (2017).
[2]A. Bohlen and T. Seefeld, "Adaptive powder nozzle setup for enhanced efficiency in laser metal deposition," J. Laser Appl. **36**, 012017 (2024).
[3]T. Schopphoven, "Experimental and model-theoretical investigations on extreme high-speed laser material deposition," Ph.D. dissertation, RWTH Aachen University, 2019.
[4]L. Zhou, G. Ma, H. Zhao, H. Mou, J. Xu, W. Wang, Z. Xing, Y. Li, W. Guo, and H. Wang, "Research status and prospect of extreme high-speed laser cladding technology," Opt. Laser Technol. **168**, 109800 (2024).
[5]V. Glushych, N. Dall, M. Zimmermann, T. Schopphoven, W. Meiners, and C. L. Häfner, "Fundamentals of simultaneous machining and coating (SMaC) through combination of extreme high-speed laser material deposition (EHLA) and turning," J. Laser Appl. **36**, 042037 (2024).
[6]C. Lampa and I. Smirnov, "High speed laser cladding of an iron based alloy developed for hard chrome replacement," J. Laser Appl. **31**, 022511 (2019).
[7]R. Garnett, *Bayesian Optimization* (Cambridge University, Cambridge, 2023).
[8]K. Wang, W. Liu, Y. Hong, H. Sohan, Y. Tong, Y. Hu, M. Zhang, J. Zhang, D. Xiang, H. Fu, and J. Ju, "An overview of technological parameter optimization in the case of laser cladding," Coatings **13**, 496 (2023).
[9]N. Pirch, S. Linnenbrink, A. Gasser, K. Wissenbach, and R. Poprawe, "Analysis of track formation during laser metal deposition," J. Laser Appl. **29**, 022506 (2017).
[10]P. Lyu, P. Li, and K. Jiang, "Numerical simulation and experimental investigation of the laser cladding processes of Ti6Al4V with coaxial shroud protection," J. Mater. Res. Technol. **26**, 4133–4150 (2023).
[11]F. Wirth, "Process understanding, modeling and predictive simulation of laser cladding." Ph.D. thesis, ETH Zürich, 2018.

[12]T. Menold, V. Onuseit, M. Buser, M. Haas, N. Bär, and A. Michalowski, "Laser material processing optimization using Bayesian optimization: A generic tool," Light: Adv. Manuf. **5**, 1 (2024).
[13]M. Kröger, B. Bornschlegel, T. Kaster, C. Hinke, and C. Holly, "Autonomous ultra short pulse ablation process design with Bayesian optimization," Proc. SPIE **13351**, 62 (2025).
[14]B. Kavas, E. C. Balta, M. R. Tucker, R. Krishnadas, A. Rupenyan, J. Lygeros, and M. Bambach, "*In situ* controller autotuning by Bayesian optimization for closed-loop feedback control of laser powder bed fusion process," arXiv:2406.19096.
[15]T. Chepiga, P. Zhilyaev, A. Ryabov, A. P. Simonov, O. N. Dubinin, D. G. Firsov, Y. O. Kuzminova, and S. A. Evlashin, "Process parameter selection for production of stainless steel 316L using efficient multi-objective Bayesian optimization algorithm," Materials **16**, 1050 (2023).
[16]V. Karkaria, A. Goeckner, R. Zha, J. Chen, J. Zhang, Q. Zhu, J. Cao, R. X. Gao, and W. Chen, "Towards a digital twin framework in additive manufacturing: Machine learning and Bayesian optimization for time series process optimization," arXiv:2402.17718.
[17]J. Sousa, A. Sousa, F. Brueckner, L. P. Reis, and A. Reis, "Human-in-the-loop multi-objective Bayesian optimization for directed energy deposition with *in situ* monitoring," Robot. Comput.-Integr. Manuf. **92**, 102892 (2025).
[18]M. Brucki, *Extremes Hochgeschwindigkeits-Laserauftragschweißen—Prozesstechnologische Untersuchung der Oberflächenbeschaffenheit von Beschichtungen*, First Auflage [Extreme High-Speed Laser Material Deposition—Process-Based Investigation of the Surface Quality of Coatings (in English)] (Apprimus, Aachen, 2023).
[19]Q. Yan, K. Yang, Z. D. Wang, M. Z. Chen, G. F. Sun, and Z. H. Ni, "Surface roughness optimization and high-temperature wear performance of H13 coating fabricated by extreme high-speed laser cladding," Opt. Laser Technol. **149**, 107823 (2022).
[20]E. Willenborg, *Polieren von Werkzeugstählen mit Laserstrahlung*, (Shaker, Aachen, 2006).
[21]J. Gardner, M. Kusner, Zhixiang, K. Weinberger, and J. Cunningham, "Bayesian optimization with inequality constraints," in *Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 22–24 June 2014* (PMLR, 2014), Vol. 32, pp. 937–945.
[22]A. von Rohr, D. Stenger, D. Scheurenberg, and S. Trimpe, "Local Bayesian optimization for controller tuning with crash constraints," at-Automatisierungstechnik **72**, 281–292 (2024).
[23]D. Stenger and D. Abel, "Benchmark of Bayesian optimization and metaheuristics for control engineering tuning problems with crash constraints," arXiv:2211.02571 (2022).
[24]C. Lara, see https://botorch.org/docs/batching/ for "Batching | BoTorch," 2025 (last accessed 27 June 2025).
[25]S. Ament, S. Daulton, D. Eriksson, M. Balandat, E. Bakshy, "Unexpected improvements to expected improvement for Bayesian optimization," arXiv:2310.20708 (2023).