



Cite this: DOI: 10.1039/d5dd00556f

# Using Flory–Huggins-informed human-in-the-loop Bayesian optimization to map the phase diagram of polymer blends

Justin C. Hughes,<sup>a</sup> Dylan J. York,<sup>b</sup> Kevin G. Yager,<sup>c</sup> Chinedum O. Osuji<sup>b</sup> and Russell J. Composto<sup>b\*</sup>

Mapping the phase diagram of polymer blends is an essential step in controlling the structure–property relationship of polymer-based materials. However, traditional grid-based approaches are inefficient and rely on subjective judgements for terminating the experimental campaign. Artificial intelligence-guided experimentation offers a compelling alternative, especially when data-driven decision-making is interfaced with established polymer thermodynamics to improve efficiency and interpretability. Here, we introduce a physics-informed Bayesian optimization approach to guide the mapping of the phase diagram of a model blend containing poly(methyl methacrylate) and poly(styrene-*ran*-acrylonitrile). Physical information is derived from a Flory–Huggins representation of the spinodal curve, which is integrated into the Bayesian optimization process as a structured prior mean that acts as a soft constraint. Implemented as a human-in-the-loop workflow, the approach leverages optical imaging of film cloudiness with iterative Gaussian process surrogate modeling and a parameter selection decision policy to identify the composition–temperature conditions for sequential iterations. Convergence of kernel and Flory–Huggins-based hyperparameters provided a stopping criterion, ensuring an objective and interpretable termination of the experimental campaign. The framework recovered the known lower critical solution temperature ( $\sim 160$  °C), while increasing material efficiency through targeted sampling. This work establishes a proof-of-concept for the application of Bayesian optimization workflows to study polymer blend miscibility.

Received 12th December 2025  
Accepted 2nd March 2026

DOI: 10.1039/d5dd00556f

rsc.li/digitaldiscovery

## Introduction

Polymer blends are fundamentally important in both commercial and research contexts because they enable tunable material properties that cannot be achieved with single-component polymer systems. Through controlled mixing of chemically distinct polymers, blends can exhibit enhanced mechanical properties,<sup>1</sup> permeability,<sup>2</sup> and thermal stability,<sup>3</sup> making them attractive for various technologies such as food packaging<sup>4</sup> and biocompatible materials.<sup>5</sup> The ability to tune macroscopic properties for applications arises from molecular interactions that govern polymer mixing, morphology and phase behavior, which ultimately dictate the structure–property relationship of materials. Foundational theory from Flory,<sup>6</sup> de Gennes,<sup>7</sup> and others established the framework to understand the relation between these interactions and predicting polymer

thermodynamics. This improved understanding along with new experimental techniques led to a large increase in polymer blending in the 1970s and 80s.<sup>8</sup> Given this extensive commercial and research background, polymer blends provide an ideal platform for validating physics-informed decision-making frameworks that guide experimental campaigns.

The thermodynamics governing polymer miscibility can be described by the Flory–Huggins (FH) theory. The FH theory describes the phase separation behavior of polymer blends to the balance between enthalpic and entropic contributions to free energy.<sup>6</sup> Polymer blends can exhibit upper critical or lower critical solution temperature behavior (UCST and LCST), for blends that phase separate upon cooling or heating, respectively. The boundary between single-phase and phase-separated regimes is defined by a composition–temperature phase diagram. At equilibrium, annealing an AB blend in the two-phase region results in the formation of domains rich and poor in the A component. When domains become large enough (micron scale), they scatter light and appear cloudy. Therefore, visual inspection of polymer blends under an optical microscope provides a simple method for ascertaining phase separation in polymer blends.<sup>9</sup> Spectroscopy,<sup>10</sup> X-ray scattering,<sup>11</sup> neutron scattering,<sup>12</sup> and electron microscopy<sup>13</sup> methods can also be used to determine polymer

<sup>a</sup>Department of Materials Science and Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. E-mail: composito@seas.upenn.edu

<sup>b</sup>Department of Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA

<sup>c</sup>Center for Functional Nanomaterials, Brookhaven National Laboratory, Upton, NY 11973, USA



blend miscibility. Since the 1970's, FH and other models (*e.g.*, self-consistent mean-field) have combined with new experimental methods to map phase behavior of 100's of combinations of binary polymer blends, including those containing compatibilizers.<sup>8,10,14</sup> However, as new polymers continue to be discovered,<sup>15–18</sup> extensive testing of possible combinations of polymers becomes problematic. Because polymers have many molecular characteristics, such as monomer type, molecular weight, structure (linear *vs.* branched) and dispersity, this large parameter space makes traditional grid-based sampling methods for mapping phase diagrams and determining thermodynamic properties challenging.<sup>19,20</sup> Furthermore, grid-based sampling efforts also depend on subjective judgements about when sufficient data has been collected. By leveraging intelligent mechanisms for extracting information from collected data and guiding experimental decisions, AI-driven methods can efficiently enable systematic mapping of thermodynamic behavior without requiring exhaustive sampling of the entire high-dimensional parameter space.

Recent adaptations of AI approaches to polymer and material science have shown promise for accurate property mapping and optimization across different material systems. Applications range from data-driven materials discovery for clean energy technologies,<sup>21</sup> to mapping the phase diagram of self-assembled block copolymers<sup>22</sup> and polymer solutions,<sup>23</sup> or lattice models such as the Ising system,<sup>24</sup> to the design of full-stack self-driving laboratory architectures that function in a closed-loop manner to autonomously synthesize and characterize materials.<sup>25,26</sup> However, despite this progress, AI-guided experimentation has seen limited application to solid-state polymer blend miscibility, where current experimental workflows are largely manual and grid-based. Importantly, these studies indicate that AI-driven experimental design is well-positioned to be applied to solid-state polymer miscibility, even though this opportunity has yet to be explored.

One effective approach for implementing such AI-guided experimental design is through Bayesian optimization (BO), a mathematical framework widely used in optimal experimental design (OED) and autonomous experimentation (AE) to model scientific data and guide decision-making. Specifically, BO is leveraged to explore or optimize an objective function that would be expensive to map exhaustively. Examples of objective functions include miscibility of a polymer blend as a function of composition, temperature, and other experimental variables, or conductivity of a polymer nanocomposite material as a function of filler weight fraction and polymer identity. In a typical implementation, BO is used to recommend new experimental parameters that maximize an acquisition function that directs measurements toward the most valuable or informative samples.<sup>27</sup> The acquisition function depends on a surrogate model that represents the optimizer's current understanding of the objective function. The most common surrogate modeling framework for this is a Gaussian process (GP), which defines a distribution over possible functions. Using a GP enables the optimizer to model both the objective's mean and uncertainty over the parameter space.<sup>28</sup> Including the uncertainty modeling component helps promote robustness in the surrogate

modeling process, in particular for data-sparse or noisy regimes of the parameter space. Many software packages for performing this surrogate modeling and point selection have been previously developed,<sup>29–31</sup> and are widely used across various experimental and computational applications.<sup>22,32,33</sup>

BO is a non-parametric, data-driven method that does not require prior knowledge to make decisions. However, for complex problems, it is often useful to integrate prior knowledge in some form. Several strategies exist to accomplish this task, including transforming the parameter space to allow for more robust surrogate modeling for parameters originally on different scales<sup>34</sup> and kernel selection to enforce the correct form for correlations in the parameter space.<sup>35</sup> Directly leveraging a structured prior function to encode previously-discovered physical laws or constraints has also been demonstrated as an effective way to augment the capabilities of the BO process.<sup>36</sup> Together, these approaches demonstrate different possible methods of integrating prior knowledge to improve the BO process.

In this work, we developed an integrated platform that combines Flory–Huggins-informed Bayesian optimization with human-driven experimental sample preparation and characterization. As a proof of concept, we apply this human-in-the-loop (HITL) approach to map the phase diagram of a model polymer blend system of poly(methyl methacrylate)/poly(styrene-*ran*-acrylonitrile) (PMMA/SAN),<sup>37</sup> leveraging cloud point to characterize miscibility. Flory–Huggins theory is embedded in a trainable structured prior mean to guide the mapping process. In this iterative workflow, each round of optimization involved preparing and annealing five polymer blend samples chosen by the Bayesian optimizer to maximize the information gained from each sample. This division of tasks, with algorithmic decision-making by the optimizer and human execution of experiments, allowed the workflow to operate without the need for investment in robotic automation. In this context, the role of the HITL approach is to enable practical deployment, while algorithmic efficiency is derived from the underlying physics-informed BO framework. Successive iterations refined the surrogate model until convergence was reached, as indicated by stabilization of hyperparameters. Using explicit convergence criteria helped reduce subjective decision-making regarding when to terminate the experimental campaign. Through this methodology, we identify both the advantages and challenges of leveraging physics-informed HITL experimentation and highlight the potential of this framework to enable efficient, interpretable study of more complex polymer systems.

## Materials and methods

### Materials

Methyl isobutyl ketone (MIBK, ACS-grade), chloroform (HPLC-grade), and methanol (HPLC-grade) were all purchased from Fisher Chemical and were used as-received. The poly(methyl methacrylate) (PMMA) used has a weight-average molecular weight ( $M_w$ ) of 84.5 kg mol<sup>-1</sup> and dispersity ( $M_w/M_n$ ) of 1.05. The PMMA was purchased from Polymer Source Inc. and was used as-received. The poly(styrene-*ran*-acrylonitrile) (SAN) used had



a weight-average molecular weight of  $118 \text{ kg mol}^{-1}$ , a dispersity of 2.24, and was 33 wt% acrylonitrile. The SAN was provided by Monsanto and was used after purification. A 3 wt% SAN solution in chloroform was added to methanol at a 1:10 volume ratio. After precipitation for 1 hour, the solvent was removed, and the precipitate was collected and dried for 48 hours. The SAN precipitate was then re-dissolved in chloroform, and the cycle was repeated two more times, for a total of three purification cycles. Plain glass microscope slides 3 in. by 1 in. by 1 mm were purchased from Fisher Scientific. Single-side-polished, phosphorus-doped silicon wafers ( $100 \pm 0.5 \text{ mm}$  diameter,  $500 \pm 25 \mu\text{m}$  thickness, 1–10  $\Omega \text{ cm}$  resistivity) were purchased from University Wafer.

### Sample preparation

Separate 5 wt% solutions of PMMA and SAN in MIBK were prepared to serve as stock solutions for the experimental campaign. Each stock solution was stirred overnight at 600 rpm. For a given experiment, 60  $\mu\text{L}$  of sample solutions were prepared in 2 mL glass vials and stirred at 300 rpm for 2 hours. Then two 20  $\mu\text{L}$  aliquots of each sample solution selected at a given iteration were drop-cast onto a glass slide. The glass slides were annealed on a hot plate at 120  $^{\circ}\text{C}$  for 45 minutes to evaporate the MIBK solvent. Subsequent annealing was performed in a Shel Lab SVAC1 series oven to promote phase separation for temperatures ranging from 150 to 200  $^{\circ}\text{C}$ . Annealing times were calibrated by measuring the time required for a series of 50:50 (wt%) PMMA/SAN samples to reach saturation in cloudiness at various temperatures above the lower critical solution temperature (LCST). These data were fit to an Arrhenius relation, which then was used to compute appropriate annealing times to accommodate a range of five compositions to be annealed simultaneously. Additional details of the calibration procedure are provided in the SI.

The parameter space for the experimental campaign was bounded by compositions ranging from 0 to 100 wt% PMMA and post-solvent evaporation temperatures ranging from 150  $^{\circ}\text{C}$  to 200  $^{\circ}\text{C}$ . These boundaries were chosen to mimic the experimental conditions reported by Newby *et al.*<sup>37</sup> (0 to 100 wt% PMMA, 150 to 200  $^{\circ}\text{C}$ ), ensuring comparability of the results from experiments. Four initial samples were prepared to seed the experimental campaign: 0 wt% PMMA at 150  $^{\circ}\text{C}$ , 100 wt% PMMA at 150  $^{\circ}\text{C}$ , 0 wt% PMMA at 200  $^{\circ}\text{C}$ , and 100 wt% PMMA at 200  $^{\circ}\text{C}$ . After this initial round of measurements was conducted, each subsequent iteration involved preparing five additional compositions at a given temperature, selected using the Bayesian optimization approach described below. This batch size (five compositions) was based on the oven's throughput limit of ten samples per annealing run: selecting five compositions per iteration allowed for duplicate samples at each composition, balancing the need for reproducibility with efficient exploration of the composition space.

### Cloud point imaging and data analysis

Films were imaged using an Angetube 862Pro digital camera (26 mm focal length, f/1.5 aperture) and a JH Technologies

LR92240 Halogen Light Source. During measurement, the glass slides were placed onto a silicon wafer and a custom-milled imaging box. After the box was closed, the halogen light source illuminated the sample from the side, enabling measurement of cloudiness. A schematic of the imaging setup is shown in SI (Fig. S1).

Images were acquired using the OpenCV library in Python. Each image was converted to grayscale using the standard luminance formulation as defined in ITU-R BT.601-7.<sup>38</sup> The formula for this conversion from a three-channel BGR image to a single-channel grayscale image is given in eqn (1).  $I_{\text{Red}}$ ,  $I_{\text{Green}}$ , and  $I_{\text{Blue}}$  denote the individual channel contributions at a given pixel, and  $I$  is the resulting grayscale pixel intensity.

$$I = 0.299 \times I_{\text{Red}} + 0.587 \times I_{\text{Green}} + 0.114 \times I_{\text{Blue}} \quad (1)$$

Afterwards, the average pixel intensity of the sample region ( $I_{\text{S}}$ ) was compared to the background intensity ( $I_{\text{B}}$ ). These values were combined with the sample-to-light source distance ( $l_{\text{S}}$ ) to compute a cloudiness score according to eqn (2). In the grayscale images, pixel intensity values ranged from 0 (black) to 255 (white), inclusive.

$$\text{Cloudiness} = \frac{1}{2} [1 + \tanh(0.015 \times ((I_{\text{S}} - I_{\text{B}}) \times l_{\text{S}}^{0.225}) - 80)] \quad (2)$$

Due to smooth variation of the hyperbolic tangent between  $-1$  and  $+1$ , along with additional scaling and shifting, the cloudiness score is normalized between 0 and 1, with the phase boundary defined by the intermediate value of 0.5. This cloudiness value is then used as the corresponding measurement for each individual sample in the Bayesian optimization process. Calibration details for the constants used in eqn (2) are provided in the SI.

### Bayesian optimization via Gaussian process surrogate modeling

Bayesian optimization (BO) was implemented using gpCAM (version 8.1.9),<sup>39</sup> with auxiliary Python scripts for imaging and analysis provided in the Zenodo repository. At each iteration, cloudiness scores were appended to the dataset, and a Gaussian process surrogate was updated with globally optimized hyperparameters using differential evolution. The structured prior mean  $\mu_0$  for this process incorporated Flory–Huggins parameters  $A$  and  $B$  to encode the spinodal construction of the phase boundary, namely where the second derivative of the free energy is zero. In this formulation, the Flory–Huggins interaction parameter  $\chi$  is defined in eqn (3), where  $A$  and  $B$  represent the entropic and enthalpic contributions, respectively, to the thermodynamics of mixing two dissimilar polymers. The interaction parameter  $\chi$  is dimensionless, with  $B$  expressed in units of Kelvin so that  $B/T$  is dimensionless, consistent with standard Flory–Huggins definitions.

$$\chi = A + \frac{B}{T} \quad (3)$$



The sign of  $B$  determines the critical behavior as a function of temperature, namely  $B < 0$  describes a LCST blend whereas  $B > 0$  yields a UCST blend. This can be substituted into the standard expression of the second derivative of free energy given in eqn (4).

$$\frac{\partial^2 \Delta \bar{F}_{\text{mix}}}{\partial x_{\text{PMMA}}^2} = k_B T \left[ \frac{1}{N_{\text{PMMA}} x_{\text{PMMA}}} + \frac{1}{N_{\text{SAN}} (1 - x_{\text{PMMA}})} - 2\chi \right] \quad (4)$$

Lastly, this second derivative was scaled using a sigmoid function defined in eqn (5). The scaling constant  $c$  is equal to  $-2 \times 10^{23}$ . This constant only controls the sharpness of the transition across the spinodal boundary and does not affect the location of the boundary itself, which is defined by the intermediate cloudiness value of 0.5. The output of the sigmoid scaling is used as the prior mean function for the surrogate modeling.

$$\mu_0 = \frac{1}{1 + \exp\left(-c \frac{\partial^2 \Delta \bar{F}_{\text{mix}}}{\partial x_{\text{PMMA}}^2}\right)} \quad (5)$$

The sigmoid normalizes the second derivative to values between 0 and 1, placing it on the same scale as the cloudiness score, which is also normalized between 0 and 1. Because the second derivative is negative in the unstable region, the prior mean approaches 1 when it predicts phase separation and approaches 0 when it predicts miscibility.

A Matérn-3/2 kernel was used to balance the smoothness of the radial basis function (RBF) and noise introduced in experiments.  $\sigma_f^2$  is the noise variance,  $l_x$  and  $l_T$  are the characteristic length scales associated with the composition and temperature dimensions respectively, and  $r$  is the scaled interpoint distance in composition-temperature space.  $\mathbf{x} = (x, T)$  and  $\mathbf{x}' = (x', T')$  denote two points in the phase space. The kernel and scaled distance formulas are displayed below in eqn (6) and (7) respectively.

$$K(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \left(1 + \sqrt{3} r(\mathbf{x}, \mathbf{x}')\right) \exp\left(-\sqrt{3} r(\mathbf{x}, \mathbf{x}')\right) \quad (6)$$

$$r(\mathbf{x}, \mathbf{x}') = \sqrt{\left(\frac{x - x'}{l_x}\right)^2 + \left(\frac{T - T'}{l_T}\right)^2} \quad (7)$$

Together, these hyperparameters from the prior mean ( $A, B$ ) and kernel ( $\sigma_f^2, l_x, l_T$ ) constitute the five hyperparameters globally optimized *via* differential evolution at each step of the experiment. During optimization, parameters are internally linearly scaled within the defined hyperparameter bounds during mutation and recombination but are otherwise unscaled and used in their presented units during the calculation of the posterior mean.<sup>40</sup> Temperature is specified in degrees Celsius in the surrogate model and converted automatically to Kelvin for evaluation of eqn (3) and (4).

The kernel formulation in eqn (6) is directly leveraged to compute the posterior mean  $\mu$  as shown in eqn (8). In this equation,  $\mathbf{x}^*$  is the test point,  $\mathbf{X}$  is a list of all previously measured points, and  $\mathbf{y}$  is the list of computed cloudiness values for each of the previously measured points. Eqn (9)

collects the terms being added together and inverted and renames as the “total” covariance matrix  $K_{\text{tot}}$ , where the diagonal additive term is a heteroscedastic noise covariance term proportional to the squared measurement value ( $\mathbf{y}_i^2$ ) for each point.

$$\mu(\mathbf{x}^*) = \mu_0(\mathbf{x}^*) + K(\mathbf{x}^*, \mathbf{X})(K_{\text{tot}}(\mathbf{X}, \mathbf{X}))^{-1}(\mathbf{y} - \mu_0(\mathbf{X})) \quad (8)$$

$$K_{\text{tot}}(\mathbf{X}, \mathbf{X}) = K(\mathbf{X}, \mathbf{X}) + \text{diag}(0.05 \times \mathbf{y}^2) \quad (9)$$

Details on the calibration of the 0.05 scaling factor in the covariance term of eqn (9), derived from a replicate-based robustness analysis of the cloudiness metric across multiple samples, are provided in the SI.

Because computation of the posterior mean relies on matrix inversion of  $K_{\text{tot}}(\mathbf{X}, \mathbf{X})$ , the stability of this computation depends heavily on the condition of this matrix. A well-conditioned kernel ensures this computation is numerically stable, and that small perturbations or errors in the data don't disproportionately impact the posterior mean. Thus, the condition number  $\kappa$  of  $K_{\text{tot}}(\mathbf{X}, \mathbf{X})$  will also be tracked over the experimental campaign, with a large condition number indicating an ill-conditioned (numerically unstable) matrix, while a small condition number (approaching 1) indicating a well-conditioned (numerically stable) matrix. The computation of the condition number  $\kappa$  is outlined in eqn (10), where  $\lambda_{\text{max}}$  and  $\lambda_{\text{min}}$  are the maximum and minimum eigenvalues of  $K_{\text{tot}}(\mathbf{X}, \mathbf{X})$  respectively.

$$\kappa = \frac{\lambda_{\text{max}}(K_{\text{tot}}(\mathbf{X}, \mathbf{X}))}{\lambda_{\text{min}}(K_{\text{tot}}(\mathbf{X}, \mathbf{X}))} \quad (10)$$

After sample preparation, imaging, and hyperparameter retraining at a given iteration, the acquisition function  $f_a$  and decision policy leverage the posterior mean  $\mu$  and posterior covariance  $\sigma$  to select the next temperature and compositions. The acquisition function is shown in eqn (11) and includes the covariance term to encourage exploration and the target value term  $P$  (shown in eqn (12)) to encourage exploitation of useful areas of the phase space, such as those close to the phase boundary. The target term is centered at 0.5, the midpoint on the cloudiness scale between the phase-separated (cloudiness > 0.5) and single-phase (cloudiness < 0.5) regimes.

$$f_a(\mathbf{x}) = \sigma(\mathbf{x}) + 0.25 \times P(\mathbf{x}) \quad (11)$$

$$P(x) = \tanh\left(\frac{0.1}{\mu(x) - 0.5}\right) \quad (12)$$

Conceptually, this acquisition function is adapted from an Upper Confidence Bound (UCB) strategy<sup>41</sup> but replaces the direct use of the posterior mean with a task-aligned transformation  $P(x)$  that promotes sampling near the phase boundary.

The decision policy for point selection first evaluates the acquisition function across compositions at each temperature using a normalized exclusion mask, where composition and



temperature are scaled by  $\Delta x = 0.025$  wt frac. PMMA and  $\Delta T = 1$  °C respectively and requires that candidate points must lie outside a unit-radius exclusion region. It then selects the temperature with the highest average masked acquisition value. At this fixed temperature, compositions are ranked by their masked acquisition values, and five points are chosen sequentially: the highest-valued point is selected first, followed by additional points in descending masked acquisition value order until the decision policy has generated a list of five valid compositions at the given temperature.

During early development, the  $k_B T$  prefactor in the Flory–Huggins expression for the second derivative of free energy (eqn (4)) was evaluated using incorrect temperature units, introducing only a magnitude scaling of the prior without altering its qualitative structure. To verify this didn't impact the evolution of the posterior mean, the posterior mean fields from the corrected and uncorrected implementations were compared, and both workflows were found to converge to statistically indistinguishable results by the final iteration. All figures and analysis therefore use the original workflow, with the full comparison provided in the SI (Fig. S2).

During the experimental campaign, there are two main methods used for ascertaining whether convergence has been met or not upon collection of data. The first involves phase boundary location monitoring: if the location of the predicted phase boundary doesn't change, the process has converged. The second involves hyperparameter monitoring: if the hyperparameters (kernel:  $\sigma_f^2$ ,  $l_x$ ,  $l_T$ ; structured prior mean:  $A$ ,  $B$ ) remain constant, then the process has converged. If either of these criteria are met, then the experimental campaign is considered converged, and additional experiments are not needed to identify the phase diagram.

## Results and discussion

The HITL approach was tested by mapping the phase diagram of a model polymer blend of PMMA and SAN. The evolution of the experimental campaign is followed to demonstrate how the BO framework guided sampling, influenced model convergence, and outperformed a traditional grid-based approach. Simulated benchmarking was used to attribute these efficiency gains to the Flory–Huggins-informed mean formulation adopted in this work, with benchmarking results shown in the SI (Simulated Benchmarking of Flory–Huggins-Informed Mean and Kernel Approaches). First, we analyze how the decision policy shaped emergent sampling behavior over the course of the campaign. Next, we examine the convergence criterion, particularly the stabilization of hyperparameters, that governed the campaigns termination. Finally, we assess the efficacy of the BO-driven design relative to a reference study that employed a traditional grid-based sampling approach,<sup>37</sup> with a focus on sampling efficiency and interpretability of results. Over the experimental campaign, 104 total samples were prepared including 4 initialization samples and 20 iterations of 5 samples. Monitoring the evolution of the phase boundary provided valuable insight into the state of the model. However, convergence was formally determined by stabilization

of hyperparameters, which served as the criterion for terminating the campaign.

### Phase separation measurements and model evolution

The phase state of polymer blends was scored and mapped onto a phase diagram. The raw data points, color-coded by their associated cloudiness score, as well as the temperatures at which measurements were collected, are displayed in Fig. 1. Fig. 1a displays the data points, with the phase boundary defined by the transition between cloudy samples in the phase-separated region above the LCST in yellow and the clear samples in the single-phase region below the LCST in blue. The solid line traces the location of this boundary, revealing an LCST at about 160 °C, which aligns well with literature that reported an LCST of 160 °C for a similar PMMA/SAN blend.<sup>37</sup> A visual comparison to the literature phase diagram<sup>37</sup> is provided in the SI. Three characteristic sampling regions are identified in Fig. 1a: sparsely-sampled near the outer edges of the phase space limits, the interior of the phase-separated region, and a densely-sampled zone along the phase boundary. Fig. 1b shows the annealing temperature trajectory over the 20 iterations. Note that the initial samples of pure PMMA and SAN annealed at 150 °C and 200 °C are not represented in Fig. 1b. From an initial temperature of 175 °C, the recommended temperatures decrease monotonically until iteration 3, then

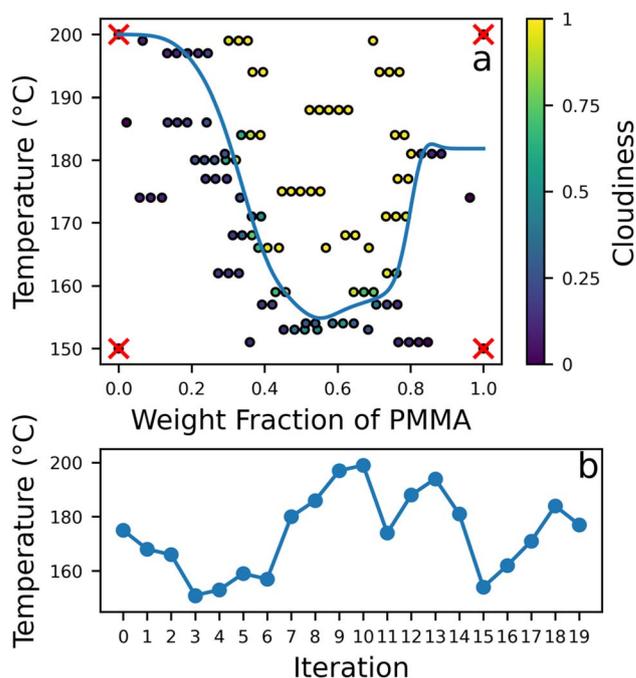


Fig. 1 Phase diagram mapping by iterations. (a) shows  $(x, T)$  data points, where yellow and blue indicate phase-separated and single-phase samples, respectively. The four initialization points are pure PMMA and SAN (red X's). The solid blue line marks the phase boundary; its flattening at low and high PMMA content arises from edge effects in the boundary-smoothing step. (b) shows the trajectory of the annealing temperatures for each set of 5 polymer blend compositions at each iteration.



increase again and vary across nearly the entire allowable temperature range.

As shown in the labeled regions in Fig. 1a, the optimizer prioritized informative measurements and avoided measurements that would provide little additional information. Near the compositional limits, the optimizer was able to quickly infer the miscibility and therefore assigned few measurements. Similarly, once the optimizer roughly mapped out the location of the phase-separated region, it devoted little sampling to its interior. In contrast, the optimizer often concentrated measurements near the phase boundary since those measurements were most helpful in refining the phase boundary location. This non-uniform sampling pattern indicates that the optimizer effectively targeted measurements that maximized information gain toward the objective of mapping the phase boundary.

The temperature trajectory reveals an emergent strategy arising indirectly from the decision policy. Early in the campaign (iterations 0–6), the recommended temperatures decrease because the boundary is flatter near its minimum, allowing the fixed-temperature set of candidate compositions to lie closer to the predicted boundary and yield higher average acquisition values. These recommendations sit slightly above the true minimum because the boundary's curvature makes temperatures just above it produce marginally larger average acquisition than the minimum itself. As data accumulate and the inferred minimum shifts downward, the recommendations track this movement while remaining slightly above the updated estimate. Once enough measurements have concentrated in this region, the proximity-based exclusion removes most nearby candidates, reducing the advantage of further sampling at lower temperatures. Consequently, from iteration 7 onward, the recommended temperatures span much of the allowable range.

To analyze the behavior of the decision policy in greater detail, particularly the emergent sampling bias toward lower temperatures at early iterations, we examined the variation of the posterior mean and acquisition function over the course of

the experimental campaign. Fig. 2 displays the variation of these two functions after taking measurements at initialization and iterations 4, 9, 14, and 19. Each column is a pair of graphs representing the model state at a given iteration. The top row is the posterior mean, determined by the measurements from samples displayed as circles, and the bottom row is the acquisition function, which includes the next set of recommended samples displayed as red dots. For the first 9 iterations, represented in the first three columns, the posterior takes on either a relatively flat or non-physical representation of miscibility. However, the posterior mean graph at iteration 14 (Fig. 2d) and onwards display a more realistic shape for a polymer blend phase boundary, closely aligning a parabolic shape expected from the Flory–Huggins theory<sup>6</sup> as more data is collected. In parallel, the acquisition function consistently highlights and targets this boundary, with sharp peaks corresponding to regions of the phase space where the posterior mean takes intermediate values (green regions around 0.5) between the phase-separated ( $\sim 1$ ) and single-phase ( $\sim 0$ ) regimes.

The model state after training at iteration 4 is particularly illustrative. The acquisition surface shows that when the computed phase boundary becomes nearly horizontal over a given composition range, the decision policy preferentially selects lower temperatures, reaffirming the trend observed in the temperature trajectory in Fig. 1b. This effect appears intermittently in later iterations with relatively flat boundary regions near the minimum temperature, such as at iteration 14 (Fig. 2i). However, as measurement density increases in the lower-temperature region and the posterior mean becomes distorted through artifacts such as the non-physical double lines at low PMMA content after iteration 9 (Fig. 2c) or the irregular surface at iteration 19 (Fig. 2j), the optimizer becomes less biased toward lower temperatures, leading to more dispersed temperature recommendations.

The optimizer's tendency to recommend measurements not just near the LCST, but near the phase boundary in general, is

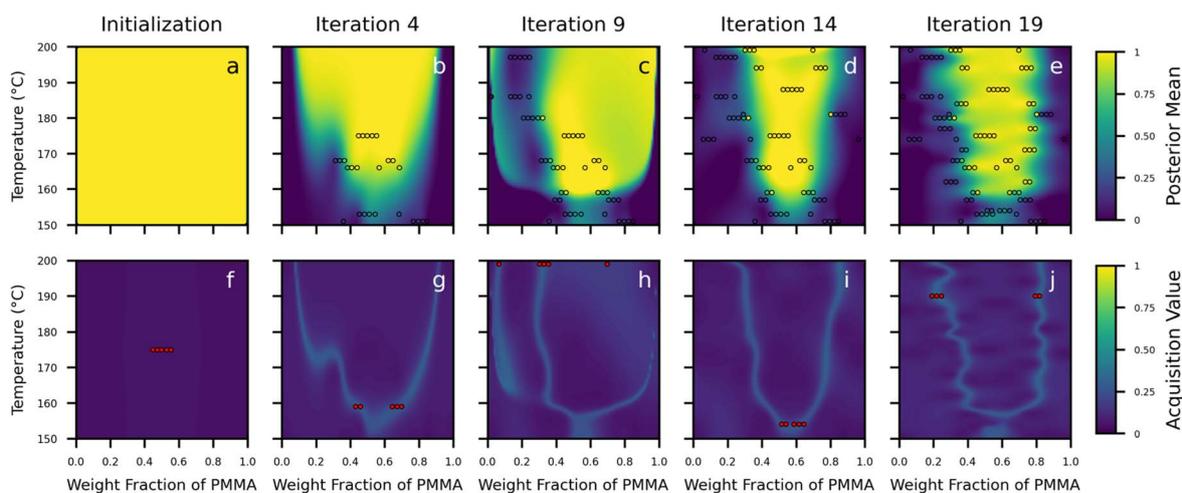
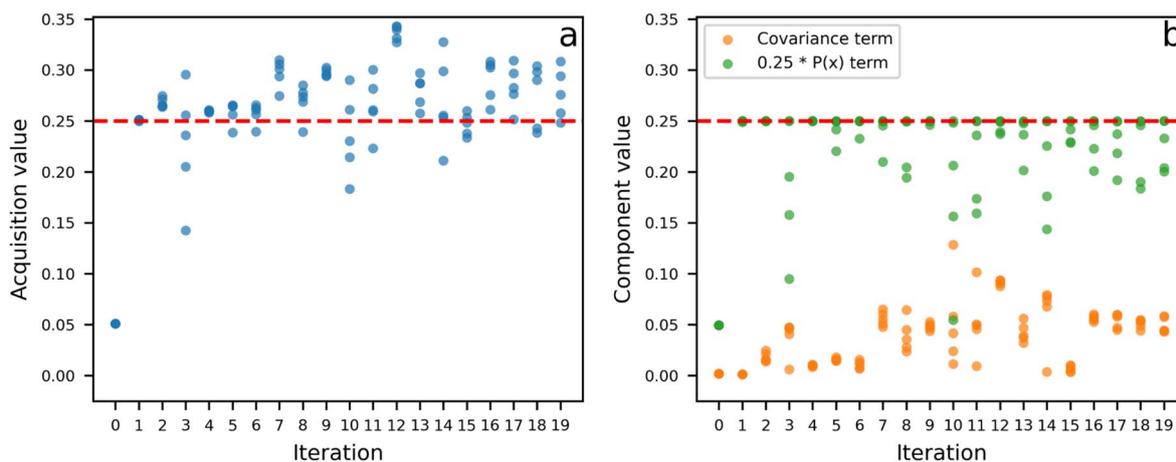


Fig. 2 Posterior mean (a–e) and acquisition function (f–j) of experimental campaign for initial state and iterations 4, 9, 14, and 19. The iterations are labeled as the iteration after which they were trained. From left to right, the graphs correspond to model states after initialization, after iteration 4, after iteration 9, after iteration 14, and after iteration 19 (the final iteration of the campaign). The red points in (f–j) represent the decision policy's recommendation for the next 5 compositions to map out the phase diagram.





**Fig. 3** Acquisition values for all 5 measured samples at a given iteration of the experimental campaign. The total acquisition function evaluation for each sample is shown in (a), and the two component terms of the acquisition function for each sample are plotted in (b). The horizontal dashed line at a value of 0.25 is included for reference for both graphs.

further reflected in the acquisition values computed from eqn (11). These acquisition values for each recommended point, evaluated at the model state when the recommendation was made, are plotted in Fig. 3. The total acquisition values plotted in Fig. 3a start at 0.05 but quickly reach and exceed values of 0.25 on average for the remainder of the campaign. Decomposing into the covariance and targeting ( $0.25 \times P(x)$ ) terms in Fig. 3b shows that the covariance term starts from around 0 but increases on average to approximately 0.05 by iteration 7. On the other hand, after iteration 0, the targeting term for most of the points selected at a given iteration is at or close to 0.25.

The decomposition shows that, by construction, the targeting term sets the dominant scale of the acquisition function and consistently favors points close to the phase boundary at each iteration. However, within this subset of candidates with similar targeting-term values, the covariance term provides the variation that determines which specific points are ultimately selected. This balance produces a non-competitive but complementary interaction between the two terms: the targeting term defines the relevant region of the parameter space, and the covariance term selects among those candidates to guide the measurements. Unlike the usual exploration–exploitation tradeoff that collapses toward a single optimum, targeting a boundary region elevates a broader set of candidates and still permits meaningful exploration among them. Together, these interactions enable high-quality point selection that supports efficient convergence toward a robust surrogate model.

### Convergence criterion monitoring

The final decision about terminating the experimental campaign was informed by monitoring the convergence criteria in terms of the change in the phase boundary location and hyperparameter convergence. The evolution of the phase boundary location over select experimental iterations, as well as the phase boundary location change over all iterations, is plotted in Fig. 4. In Fig. 4a, the boundaries at key iterations were computed using the posterior mean and the intermediate value

formulation of the phase boundary as defined in Bayesian Optimization *via* Gaussian Process Surrogate Modeling. No phase boundary can be drawn from the 4 initial samples at the corner of the phase space, as the posterior is nearly flat across the entire space (Fig. 2a), but an LCST of about 160 °C is found by iteration 4. Fig. 4b visualizes the phase boundary location change over experimental iterations, largely decreasing as the experiment progresses. There are 4 discrete regions sections of the graph: Section 1 indicates increasingly large changes as the optimizer searches for the LCST, Sections 2 and 4 indicate the optimizer refining the exact shape, leading to gradually convergent locations of the phase boundary, and Section 3 indicates a brief numerical stability influence on the exact phase boundary location.

The two graphs in Fig. 4 present a few key observations about the course of the experiment. Examining the phase boundaries directly in Fig. 4a, the general shape becomes more realistic, or similar to previous theoretical predictions about phase boundary shapes as experiments progress. The first boundary, computed from the state of the posterior after just the initialization samples, is undefined since the posterior mean did not yield a defined boundary between the two phase and single-phase regions. However, as more data accumulated in the campaign, the phase boundary reached a physically reasonable shape with a defined LCST by later iterations.

Examining the phase boundary location change plot of Fig. 4b provides a more detailed explanation of the progression of the experiment. At first, in earlier iterations of the experimental campaign (*i.e.* Section 1), large increases in the phase boundary change were observed from iterations 1 to 3. However, a near-monotonic decrease in the phase boundary location change was observed in Sections 2 and 4, meaning that the optimizer's decision policy aided in the goal of progressively refining the phase boundary location. In practice, the phase boundary location change failed as a standalone convergence metric due to the lack of convergence and the susceptibility to local minima in the hyperparameter optimization. While the



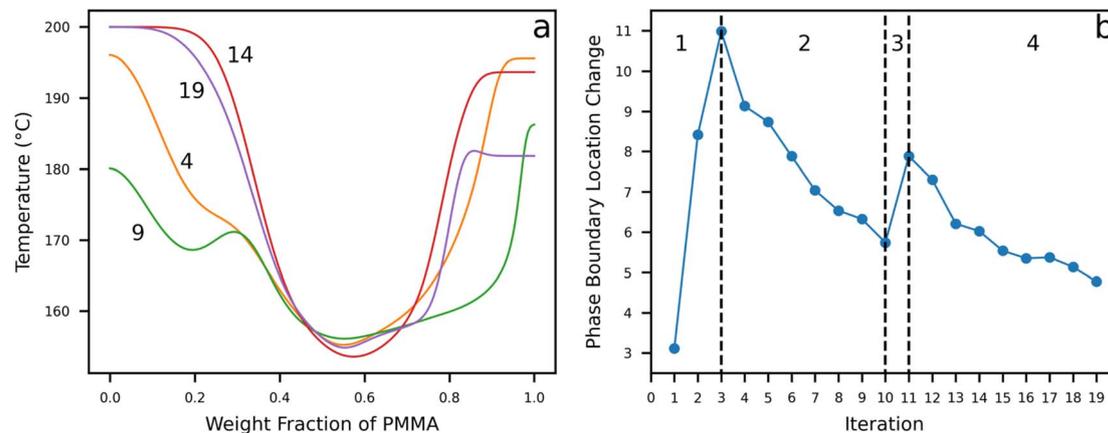


Fig. 4 Phase boundary evolution over select experimental iterations (a) and iterative boundary location change over entire experimental campaign (b). The horizontal axis on the boundary location change plot denotes the end index for subtraction, *i.e.* for 20 iterations and an initialization step, there will be 20 entries.

iterative changes are decreasing as the experiment progresses, the phase boundary location change graph doesn't appear to approach a convergence limit, even while the boundaries are taking on intuitive shapes and locations. A separate issue is the emergence of unintuitive trends, such as the spike between iterations 10 and 11 (Section 3), which can be better diagnosed by examining the variation of hyperparameters over the course of the experimental campaign.

The initial values for each of the hyperparameters, as well as their associated ranges, are shown in Table 1 below. Wide boundaries for each hyperparameter were chosen to demonstrate that the optimizer could identify reasonable fits even under poorly constrained conditions. However, further constraints can be imposed on the sign and magnitudes of  $A$  and  $B$  if the blend is known to exhibit UCST or LCST behavior.

The hyperparameters were varied iteratively leveraging the gpCAM-based implementation of differential evolution for hyperparameter fitting, with a round of training completed after the initialization and each subsequent experimental iteration. The hyperparameter logs as a function of iterations is shown in Fig. 5.  $\sigma_f^2$  values fluctuate early in the experimental campaign and eventually settles near 0.1, but this reflects the optimizer's numerical preference for low noise rather than physically interpretable behavior. For that reason, the noise variance trajectory is shown in the SI (Fig. S3). The two length scale hyperparameters  $l_x$  and  $l_T$ , shown in Fig. 5a and b respectively, initially reside around their upper bounds, gradually increasing, and then begin to converge starting at iteration 15. The final value for the composition length scale was 0.12 wt frac. PMMA and the temperature length scale was 2.63 °C. The two

Flory–Huggins parameters are initially far from the bounds, but  $B$ , as displayed in Fig. 5d, decreased until approaching the lower bound. Only after a few iterations did  $B$  converge to an intermediate value of  $-8.22$  K. Conversely,  $A$ , as displayed in Fig. 5c, initially decreased for a few iterations, increased until reaching a plateau at the same iterations that  $B$  was stuck at the lower bound, and then converged to a final value of 0.022.

While the noise variance had little effect on the model, the behavior of  $l_x$ ,  $l_T$ ,  $A$ , and  $B$  as iterations increased was informative. The two length scale hyperparameters  $l_x$  and  $l_T$  varied in a correlated manner because they are tightly coupled in the Matérn-3/2 kernel's computation of  $\mathbf{r}$ , the axis-normalized distance between two points  $\mathbf{x} = (x, T)$ ,  $\mathbf{x}' = (x', T')$  in the phase space (see eqn (7)). The Flory–Huggins parameters varied in a connected fashion because they are also tightly coupled together in computing characteristics of the phase diagram, such as the lower critical solution temperature (LCST). In particular, the variation of  $A$  and  $B$  exhibited a nearly inverse relationship. Namely as  $B(A)$  approached lower values,  $A(B)$  tended toward higher values. This relationship is consistent with their contribution to the critical temperature relation in eqn (13), which is a rearrangement of eqn (3) for  $T = T_c$ .

$$T_c = \frac{B}{\chi_C - A} \quad (13)$$

Although this relationship was implicitly hard-coded through the structured prior mean in eqn (3)–(5), it is encouraging that the variation reappeared as intended during the campaign and was not obscured by the variation of

Table 1 Initial values and boundaries for kernel and structured prior hyperparameters

Parameter	Noise ( $\sigma_f^2$ )	Composition length scale ( $l_x$ )	Temperature length scale ( $l_T$ )	Flory–Huggins parameter A	Flory–Huggins parameter B
Initial value	0.1	0.01	1	0	0
Lower bound	0.01	0.001	0.1	-1	-200
Upper bound	10	1	100	1	200



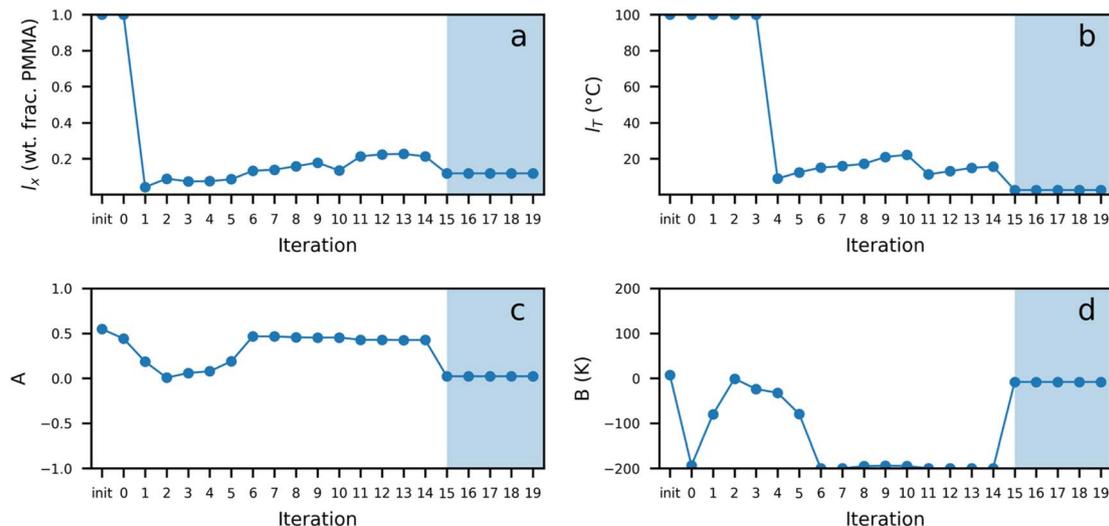


Fig. 5 Hyperparameter logs to visualize variation of kernel hyperparameters (a)  $l_x$ , (b)  $l_T$  and structured prior hyperparameters (c)  $A$ , (d)  $B$ . Vertical axis limits match the allowable ranges defined in Table 1.

hyperparameter values in the physics-agnostic kernel formulation. Additional discussion of the self-consistency of this approach is provided in the SI.

The physics-informed method's effectiveness was also evident in the evolution of the sign of  $B$  over the course of the experimental campaign. Although the PMMA/SAN blend is known to have an LCST of around 160 °C,<sup>37</sup> the optimizer was initialized without that knowledge and permitted to explore both positive and negative values for  $B$  (see Table 1). This setup forced the optimizer to implicitly determine the thermodynamic character of the blend, *i.e.* decide on LCST vs. UCST, based solely on experimental feedback and influence from the structured prior mean. After the first iterative set of samples was prepared, characterized, and their data integrated into the BO workflow, the optimizer correctly identified that  $B$  was negative, and therefore the PMMA/SAN blend exhibited an LCST. This demonstrates that the BO process can autonomously infer the correct concavity of the phase diagram without explicit external guidance.

To assess the consistency between the LCST implied by the structured prior and that extracted directly from the posterior boundary, the two values are compared in Fig. 6. The LCST was directly computed from the fitted  $A$  and  $B$  values from the structured prior mean (blue) and by extracting the minimum of the phase boundary from the posterior mean (orange). Both estimates stabilize near the expected LCST of 160 °C at iteration 5, with some deviation from this relatively constant representation between iterations 10 and 14. The inset of Fig. 6 displays the iterative differences between the two LCST representations, with the extracted value subtracted from the computed value. After iteration 2, the differences between the two representations are small and relatively constant, on the order of 2 to 4 °C, except for the same block of Iterations 10 to 14 mentioned before. In summary, both LCST estimates remain closely

aligned overall, with the cause of their intermittent deviations explored next.

Prior to iteration 4, the posterior mean fails to accurately capture the phase boundary, resulting in substantial discrepancies between the LCST values predicted by the Flory–Huggins prior and those extracted from the surrogate model. Once the phase boundary begins to learn general shapes, the structured prior and resulting posterior show good agreement. A slight increase in the Flory–Huggins parameter  $B$  between iterations 10 and 11, as shown in Fig. 5, led to a pronounced deviation between the LCST predicted by the structured prior and that extracted from the posterior model. This disagreement can be attributed to the surrogate model fitting process being temporarily being trapped in a local minimum in the hyperparameter space. While this disagreement persists for a few iterations, the

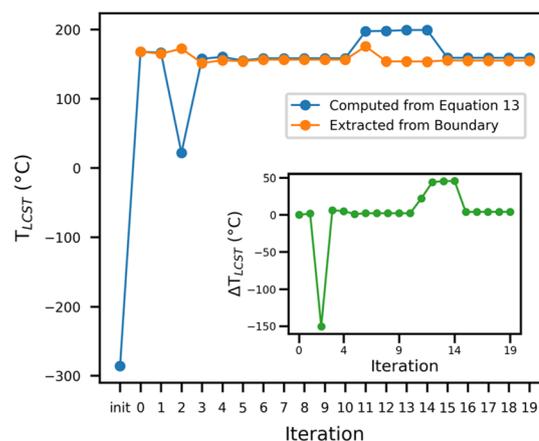


Fig. 6 Monitoring of the LCST over the experimental campaign. The LCST values from the structured prior Flory–Huggins parameters (blue) and extracted from the posterior-derived phase boundary (orange) are compared. The difference between the LCST values is plotted in the inset.



two values reconvene to similar values after iteration 15. Similar to the hyperparameter values stabilizing after this iteration, the computed and extracted LCST values reach a constant difference of about 4 °C: the final computed LCST was 158.8 °C, while the final extracted LCST was 154.8 °C. Recovering from large, near 50 °C disagreements between the structured prior and the posterior is another demonstration of the efficacy and general flexibility of this physics-informed prior formulation of the BO process.

An effective complementary approach for evaluating the numerical stability of the Bayesian optimization process involves monitoring the condition number of the kernel matrix  $K_{\text{tot}}(\mathbf{X}, \mathbf{X})$ . A plot of the evolution of the kernel matrix condition number is displayed in Fig. 7 – note the base-10 logarithmic units on the vertical axis. The kernel condition number near-monotonically increased until a peak around iteration 13. Then, at iteration 15 onward, concurrent with the iterations over which the hyperparameters converged, the kernel matrix condition converges to around 548.2. The shading on the graph displays the regime in which the condition values are converged.

The increases in condition number during early iterations indicate adjustment of hyperparameter values pertaining to the kernel as the model incorporated more data, while the stabilization past iteration 15 indicates that the kernel found a more stable regime of parameter space that better represents the phase diagram shape. Unlike the hyperparameters that stabilized values past the numerical computation limits of Python, the kernel condition numbers displayed a slight variation between iterations 15 and 19: 545.560, 548.187, 548.219, 548.219, 548.222. Since the kernel matrix values are computed leveraging the hyperparameters, the corresponding condition number remains largely unchanged once the hyperparameters stabilize. However, since matrix inversion is a global operation, a slight variation in the condition number can be directly attributed to the additional rows and columns with additional measurements. Overall, the condition number served as a good complimentary diagnostic, both capturing the impact of

hyperparameter variation as well as global effects of added measurements.

### Comparison of Bayesian optimization with a traditional phase-mapping approach

Comparing the BO framework to traditional grid-based phase mapping methods highlights how data-driven decision policies can fundamentally improve experimental efficiency and interpretability. The phase mapping study from Newby *et al.*,<sup>37</sup> which investigated the same blend of PMMA/SAN, serves as an appropriate benchmark for evaluating the BO approach against a conventional experimental design. In comparing the two approaches, three distinctions emerge: the BO framework (1) achieves greater efficiency in material usage through its decision-policy-guided sampling, (2) offers enhanced interpretability *via* physically meaningful kernel and structured prior mean hyperparameters, and (3) enables formal, quantitative definitions of convergence for increased objectivity in deciding when to terminate the experimental campaign.

A key advantage of the BO framework over traditional grid-based approaches is the promotion of material efficiency through adaptive selection of the most informative samples to be prepared and characterized. In grid-based mapping, such as that used by Newby *et al.*,<sup>37</sup> samples are often fixed at regularly-spaced compositions: for instance, PMMA weight fractions of 0.05, 0.1, 0.25, 0.5, 0.75, 0.9, and 0.95 measured at discrete temperature intervals. In contrast, the BO framework selected the combination of compositions and temperature for sample preparation to optimize the information gained from preparing each sample, without necessarily adhering to strict grid spacing. For example, the first iteration of samples post-initialization had PMMA weight fractions of 0.499, 0.527, 0.473, 0.448, and 0.553, all annealed at 175 °C, where the optimizer predicted measurements would be the most informative. The choice of these specific compositions and temperature, as well as those from subsequent iterations, allowed the BO process to focus on highly informative regions of the phase space, minimizing wasted material on experiments that wouldn't provide as much value.

Additionally, the presence of a kernel and structured prior mean enables a more informative and interpretable view of experimental progress. Because the kernel hyperparameters (the noise variance, and the two length scales  $l_x$  and  $l_T$ ) and the structured prior mean hyperparameters ( $A$  and  $B$  from the Flory–Huggins interaction parameter) have direct physical meaning, the optimization process inherently produces interpretable quantities at each iteration. This design eliminates the need for more complex code structures for *post-hoc* analysis, which represents a major advantage over traditional trial-and-error approaches, such as those used by Newby *et al.*<sup>37</sup> With this BO approach, we now have a continually-improving, theoretically-grounded understanding of the system as the experimental campaign progresses.

Lastly, the BO framework enables formal, quantitative definitions of convergence that promote transparency and reduce bias in decision-making compared to traditional approaches. In

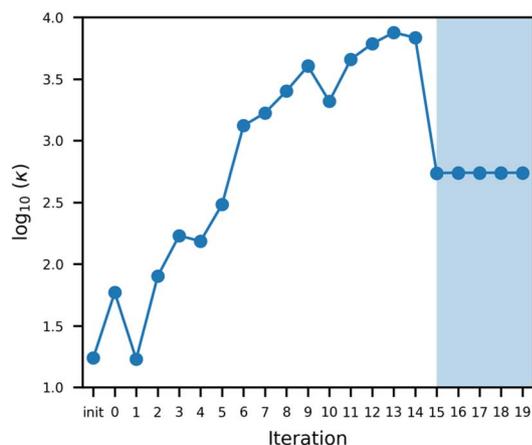


Fig. 7 Kernel matrix condition monitoring over the course of the experimental campaign. After 15 iterations, the values stabilize near 548.2 (shaded region).



traditional grid-based mapping, decisions about when to stop an experimental campaign are often defined by practical, but ill-defined, factors such as time, material availability, or intuition. The BO framework instead formalizes this process by introducing quantitative criteria that define when to stop the experimental campaign. In this work, stabilization of hyperparameter values over multiple iterations was the primary criteria, while the phase boundary location change over iterations was monitored as well. These quantitative criteria replace ad hoc decision-making with well-defined, objective rules, reducing bias across the experimental campaign. In doing so, the BO framework enhances both the rigor and interpretability of the experimental workflow.

## Conclusions

In this work, we have successfully developed a human-in-the-loop-driven experimental design framework that is able to draw suggestions from Flory–Huggins theory as to the shape of the phase diagram and use data-informed decisions to map that phase diagram. As part of this effort, we developed a unique acquisition function and decision policy that helped direct sample preparation toward regions of the phase space that would most efficiently resolve the location of the phase boundary. Explicitly monitoring the hyperparameters associated with the kernel and structured prior, as well as the iterative phase boundary shift and kernel matrix condition number, provided insight into the convergence behavior of the optimizer toward a consistent, accurate representation of the phase space. In addition to providing direct insight into the state of the fitting process, leveraging this type of monitoring was favorable because it depended on well-established mathematical constructs: previously established hyperparameters of kernels, linear algebraic concepts in the condition number, and basic algebraic operations in computing the iterative phase boundary shift. This straightforward design helps ensure that the overarching decision-making framework remains interpretable and grounded in well-defined mathematical principles.

Leveraging the customized Bayesian optimization approach, we were able to generate a phase diagram for a PMMA/SAN blend with a known experimental phase diagram. This phase diagram, while partially guided by the Flory–Huggins-based spinodal construction of the phase boundary, accurately reproduces the previously reported LCST of 160 °C,<sup>37</sup> with values extracted directly from the generated phase diagram differing by about 1 °C and those computed from the Flory–Huggins parameters by about 5 °C. The experimental campaign converged after 20 post-initialization iterations, with hyperparameters convergence defining the stopping criterion used to terminate, and the resulting posterior mean surface accurately mapped the miscibility behavior of the PMMA/SAN blend over the phase space. These results establish the capability of this human-in-the-loop approach to not only map a phase diagram but also communicate insight about physical phenomena by incorporating physical constraints through the form of the structured prior mean.

Compared to traditional grid-based mapping, the BO framework offers clear improvements in efficiency, interpretability, and objectivity. The adaptive decision policy-guided sampling helped focus the experimental campaign to target regions of the phase space that would provide the most information. Incorporating the physics-informed prior, as well as interpretable kernel structures, embedded established thermodynamic principles as well as enabled transparency during the experimental campaign. Finally, the implementation of convergence criteria in the hyperparameter stabilization and phase boundary location change convergence provided objective, rule-based decision-making that replace subjective judgement calls incumbent with traditional methods. Together, these features make the BO framework measurably better than traditional grid-based approaches for phase mapping.

Beyond demonstrating the fidelity of this decision-making framework to mapping polymer blends, this framework represents a generalizable strategy for integrating domain-specific physical constraints with Bayesian optimization in experimental science. The combination of intentionally designed structured priors and interpretable convergence metrics, combined with the relatively low barrier-to-entry of designing code constructs as opposed to implementing robotic control of experiments, presents a powerful foundation for future autonomous experimental work that requires both efficiency and interpretability, without requiring extensive robotic or computer science expertise. Crucially, this separation of algorithmic decision-making from experimental execution allows such workflows to be deployed in manual laboratory settings, expanding the practical application of physics-informed BO workflows beyond fully autonomous platforms. Future work will focus on extending this approach to higher-dimensional parameter spaces associated with polymer-based materials and helping optimize structure–property relations for applications of those systems.

An alternative extension of this framework would be to replace the Gaussian process surrogate with a Bayesian mechanistic model derived from an underlying physical theory. In such an approach, model parameters, for example  $A$  and  $B$  values that determine the Flory–Huggins interaction parameter  $\chi$ , could be fit directly from experimental observations within the iterative experimentation loop. While reducing the complexity of the computational approach to that of a lower-dimensional regression task, this strategy introduces significant trade-offs including reduced flexibility of the workflow. When the governing model is incomplete or only approximately describes the system, mismatches between theory and the experiment may compound with numerical fitting approximations to yield poor results. For this reason, future work will continue to leverage Gaussian process surrogates designed to incorporate physical knowledge as a soft constraint, such as through prior means, to allow physical insight to guide experimental execution without over-constraining the surrogate model.

The foundation established in this study provides a starting point for extending this framework beyond model systems with well-characterized phase behavior to unknown and more



complex polymer systems. In multicomponent blends, the dimensionality of the composition space, as well as the number of parameters associated with each component (polymer identity, molecular weight, dispersity, *etc.*), grows rapidly, making exhaustive mapping of thermodynamic relationships difficult. In this framework, physics-informed priors can provide early guidance toward experimental goals. However, this influence is progressively diminished as experimental data accumulates because physical knowledge is incorporated as a soft constraint rather than a hard constraint, preventing incorrect prior assumptions from constraining the final posterior. This balance between physical guidance and asymptotic data dominance enables physics-informed approaches to remain effective as the system complexity increases, even when available physical insight is limited. Beyond polymer blends, similar principles can be applied to polymer nanocomposites, where the variety of inorganic fillers, polymer matrices, and processing parameters demands a combination of computational guidance and physical intuition that a physics-informed BO framework can provide. In both classes of systems, embedded domain knowledge through a structured prior can accelerate discovery while maintaining objectivity and interpretability in the experimental process.

## Author contributions

J. C. H.: conceptualization, methodology, software, validation, formal analysis, investigation, data curation, writing – original draft, writing – review & editing, visualization. D. J. Y.: methodology, software, formal analysis, investigation. K. G. Y.: resources, writing – review & editing, supervision. C. O. O.: supervision, project administration, funding acquisition. R. J. C.: conceptualization, resources, writing – review & editing, visualization, supervision, project administration, funding acquisition.

## Conflicts of interest

The authors declare no competing financial interest.

## Abbreviations

$A$	Parameter describing the strength of entropic interactions in polymer mixtures Flory–Huggins theory (unitless)
$B$	Parameter describing the strength of enthalpic interactions in polymer mixtures according to Flory–Huggins theory (K)
$\chi$	Flory–Huggins interaction parameter (unitless)
$l_x$	Composition length scale (weight fraction (unitless))
$l_T$	Temperature length scale (°C)
$x$	Composition of individual sample (weight fraction (unitless))

$T$	Temperature of annealing for individual sample (°C)
$x$	Individual $(x, T)$ point in phase space corresponding to a sample or prospective sample (weight fraction, °C)
$X$	List of parameters corresponding to previously prepared and characterized samples (list of (weight fraction, °C))
$y$	List of cloudiness values for measured samples (list of unitless values)
Acquisition function	Function to estimate the utility of sampling a point in parameter space, balancing exploration and exploitation
Autonomous experimentation	Automated execution of experiments governed by choice of kernel, decision policy, and other functional forms
Bayesian optimization	Method for black-box function optimization relying on probabilistic surrogate modeling for decision-making
Condition number	Describes the stability of a mathematical operation with respect to changes or errors in the input data
Decision policy	Function that takes in acquisition function evaluations and recommends sample points for measurement
Differential evolution	Population-based stochastic optimization method that does not require gradient computation
Gaussian process	Stochastic method to generate distributions of functions over a parameter space, defined by a mean and covariance
Human-in-the-loop	Partially automated execution of experiments that leverage some amount of human intervention
Hyperparameters	Trainable parameters for components of a Gaussian process that impact its representation of the objective function
Optimal experimental design	Paradigm of experimental design that selects conditions to optimize a statistic, such as information gain, at each step of the experimental campaign
Posterior covariance	Model's uncertainty estimation between points, given observed data and encoded prior knowledge
Posterior mean	Model's prediction of objective function value, given observed data and encoded prior knowledge
Structured prior mean	A function encoding prior knowledge of the system, such as governing physical equations or constraints

## Data availability

The source code for the Bayesian optimization workflow, as well as the corresponding data, Jupyter Notebooks, and CAD file for the imaging setup, is available on GitHub at <https://github.com/>



[jhughes3/fh-hitl-bo](#), with an archived version on Zenodo at <https://doi.org/10.5281/zenodo.18805553>.

Supplementary information (SI) is available: annealing time and parameter calibration, details about imaging setup, code analysis, noise hyperparameter log, simulated benchmarking of Flory–Huggins prior efficacy, phase boundary and interaction parameter comparisons to results from previous literature. See DOI: <https://doi.org/10.1039/d5dd00556f>.

## Acknowledgements

This research is based on work supported by the National Science Foundation under Awards DMR-2407300 (J. C. H, D. J. Y., R. J. C.) and NRT-2152205 (J. C. H, C. O. O, R. J. C.). This research was also supported by the Vagelos Institute for Energy Science and Technology at the University of Pennsylvania (J. C. H.) and by the Penn Undergraduate Research Mentoring Program through the Center for Undergraduate Research and Fellowships (D. J. Y.). This research used resources of the Center for Functional Nanomaterials (CFN), which is a U.S. Department of Energy Office of Science User Facility, at Brookhaven National Laboratory under Contract No. DE-SC0012704 (K. G. Y.). The authors also gratefully acknowledge the use of facilities and instrumentation (JH Technologies LR92240 Halogen Light Source) supported by the Department of Materials Science and Engineering Departmental Laboratory at the University of Pennsylvania.

## References

- 1 A.-H. I. Mourad, Thermo-Mechanical Characteristics of Thermally Aged Polyethylene/Polypropylene Blends, *Mater. Des.*, 2010, **31**(2), 918–929, DOI: [10.1016/j.matdes.2009.07.031](https://doi.org/10.1016/j.matdes.2009.07.031).
- 2 P. M. Subramanian, Permeability Barriers by Controlled Morphology of Polymer Blends, *Polym. Eng. Sci.*, 1985, **25**(8), 483–487, DOI: [10.1002/pen.760250810](https://doi.org/10.1002/pen.760250810).
- 3 H. Tsuji and I. Fukui, Enhanced Thermal Stability of Poly(Lactide)s in the Melt by Enantiomeric Polymer Blending, *Polymer*, 2003, **44**(10), 2891–2896, DOI: [10.1016/S0032-3861\(03\)00175-7](https://doi.org/10.1016/S0032-3861(03)00175-7).
- 4 R. Muthuraj, M. Misra and A. K. Mohanty, Biodegradable Compatibilized Polymer Blends for Packaging Applications: A Literature Review, *J. Appl. Polym. Sci.*, 2018, **135**(24), 45726, DOI: [10.1002/app.45726](https://doi.org/10.1002/app.45726).
- 5 E. Brynda, M. Houska, S. P. Novikova and N. B. Dobrova, Polyethylene/Hydrophilic Polymer Blends for Biomedical Applications, *Biomaterials*, 1987, **8**(1), 57–60, DOI: [10.1016/0142-9612\(87\)90031-7](https://doi.org/10.1016/0142-9612(87)90031-7).
- 6 P. J. Flory, *Principles of Polymer Chemistry*, Cornell University Press, 1953.
- 7 P.-G. de Gennes, *Scaling Concepts in Polymer Physics*, Cornell university press, Ithaca (N.Y.) London, 1991.
- 8 D. Paul and S. Newman, *Polymer Blends*, Academic Press, New York, 1978, vol. 1, p. 2.
- 9 S. Krause, Polymer-Polymer Miscibility, *Pure Appl. Chem.*, 1986, **58**(12), 1553–1560, DOI: [10.1351/pac198658121553](https://doi.org/10.1351/pac198658121553).
- 10 E. Manias and L. A. Utracki, Thermodynamics of Polymer Blends, in *Polymer Blends Handbook*, ed. L. A. Utracki and C. A. Wilkie, Springer Netherlands, Dordrecht, 2014, pp. 171–289, DOI: [10.1007/978-94-007-6064-6\\_4](https://doi.org/10.1007/978-94-007-6064-6_4).
- 11 J. Sharma, *Characterization of Polymer Blends by X-Ray Scattering: SAXS and WAXS*, in *Characterization of Polymer Blends*, John Wiley & Sons, Ltd, 2014, pp. 209–236. DOI: [10.1002/9783527645602.ch06](https://doi.org/10.1002/9783527645602.ch06).
- 12 K. Hahn, B. J. Schmitt, M. Kirschev, R. G. Kirste, H. Salié and S. Schmitt-Strecker, Structure and Thermodynamics in Polymer Blends. Neutron Scattering Measurements on Blends of Poly(Methyl Methacrylate) and Poly(Styrene-Co-Acrylonitrile), *Polymer*, 1992, **33**(24), 5150–5166, DOI: [10.1016/0032-3861\(92\)90795-X](https://doi.org/10.1016/0032-3861(92)90795-X).
- 13 V. K. Berry, Characterization of Polymer Blends by Low Voltage Scanning Electron Microscopy, *Scanning*, 1988, **10**(1), 19–27, DOI: [10.1002/sca.4950100105](https://doi.org/10.1002/sca.4950100105).
- 14 E. Martuscelli; R. Palumbo and M. Kryszevski, *Polymer Blends: Processing, Morphology, and Properties*, 1st edn, Plenum Press, New York, 1980.
- 15 P. Shieh, H. V.-T. Nguyen and J. A. Johnson, Tailored Silyl Ether Monomers Enable Backbone-Degradable Polynorbornene-Based Linear, Bottlebrush and Star Copolymers through ROMP, *Nat. Chem.*, 2019, **11**(12), 1124–1132, DOI: [10.1038/s41557-019-0352-4](https://doi.org/10.1038/s41557-019-0352-4).
- 16 S. Ogawa, H. Morita, Y.-I. Hsu, H. Uyama and M. Tobisu, Controlled Degradation of Chemically Stable Poly(Aryl Ethers) via Directing Group-Assisted Catalysis, *Chem. Sci.*, 2024, **15**(42), 17556–17561, DOI: [10.1039/D4SC04147J](https://doi.org/10.1039/D4SC04147J).
- 17 M. Hassan, G. A. Bhat and D. J. Darensbourg, Post-Polymerization Functionalization of Aliphatic Polycarbonates Using Click Chemistry, *Polym. Chem.*, 2024, **15**(18), 1803–1820, DOI: [10.1039/D4PY00174E](https://doi.org/10.1039/D4PY00174E).
- 18 J. N. Lalonde, G. Pilania and B. L. Marrone, Materials Designed to Degrade: Structure, Properties, Processing, and Performance Relationships in Polyhydroxyalkanoate Biopolymers, *Polym. Chem.*, 2025, **16**(3), 235–265, DOI: [10.1039/D4PY00623B](https://doi.org/10.1039/D4PY00623B).
- 19 S. M. Moosavi, K. M. Jablonka and B. Smit, The Role of Machine Learning in the Understanding and Design of Materials, *J. Am. Chem. Soc.*, 2020, **142**(48), 20273–20287, DOI: [10.1021/jacs.0c09105](https://doi.org/10.1021/jacs.0c09105).
- 20 T. Long, Q. Pang, Y. Deng, X. Pang, Y. Zhang, R. Yang and C. Zhou, Recent Progress of Artificial Intelligence Application in Polymer Materials, *Polymers*, 2025, **17**(12), 1667, DOI: [10.3390/polym17121667](https://doi.org/10.3390/polym17121667).
- 21 D. P. Tabor, L. M. Roch, S. K. Saikin, C. Kreisbeck, D. Sheberla, J. H. Montoya, S. Dwaraknath, M. Aykol, C. Ortiz, H. Tribukait, C. Amador-Bedolla, C. J. Brabec, B. Maruyama, K. A. Persson and A. Aspuru-Guzik, Accelerating the Discovery of Materials for Clean Energy in the Era of Smart Automation, *Nat. Rev. Mater.*, 2018, **3**(5), 5–20, DOI: [10.1038/s41578-018-0005-z](https://doi.org/10.1038/s41578-018-0005-z).
- 22 G. S. Doerk, A. Stein, S. Bae, M. M. Noack, M. Fukuto and K. G. Yager, Autonomous Discovery of Emergent Morphologies in Directed Self-Assembly of Block



- Copolymer Blends, *Sci. Adv.*, 2023, 9(2), eadd3687, DOI: [10.1126/sciadv.add3687](https://doi.org/10.1126/sciadv.add3687).
- 23 J. G. Ethier, D. J. Audus, D. C. Ryan and R. A. Vaia, Integrating Theory with Machine Learning for Predicting Polymer Solution Phase Behavior, *Giant*, 2023, 15, 100171, DOI: [10.1016/j.giant.2023.100171](https://doi.org/10.1016/j.giant.2023.100171).
- 24 I. Jang and A. Yethiraj, Unsupervised Machine Learning Method for the Phase Behavior of the Constant Magnetization Ising Model in Two and Three Dimensions, *J. Phys. Chem. B*, 2025, 129(1), 532–539, DOI: [10.1021/acs.jpcc.4c06261](https://doi.org/10.1021/acs.jpcc.4c06261).
- 25 P. A. Beaucage and T. B. Martin, The Autonomous Formulation Laboratory: An Open Liquid Handling Platform for Formulation Discovery Using X-Ray and Neutron Scattering, *Chem. Mater.*, 2023, 35(3), 846–852, DOI: [10.1021/acs.chemmater.2c03118](https://doi.org/10.1021/acs.chemmater.2c03118).
- 26 A. E. Gongora, B. Xu, W. Perry, C. Okoye, P. Riley, K. G. Reyes, E. F. Morgan and K. A. Brown, A Bayesian Experimental Autonomous Researcher for Mechanical Design, *Sci. Adv.*, 2020, 6(15), eaaz1708, DOI: [10.1126/sciadv.aaz1708](https://doi.org/10.1126/sciadv.aaz1708).
- 27 B. Shahriari, K. Swersky, Z. Wang, R. P. Adams and N. de Freitas, Taking the Human Out of the Loop: A Review of Bayesian Optimization, *Proc. IEEE*, 2016, 104(1), 148–175, DOI: [10.1109/JPROC.2015.2494218](https://doi.org/10.1109/JPROC.2015.2494218).
- 28 C. K. I. Williams and C. E. Rasmussen, *Gaussian Processes for Regression*, MIT Press, Cambridge, Massachusetts, 1995, pp. 514–520.
- 29 M. M. Noack, P. H. Zwart, D. M. Ushizima, M. Fukuto, K. G. Yager, K. C. Elbert, C. B. Murray, A. Stein, G. S. Doerk, E. H. R. Tsai, R. Li, G. Freychet, M. Zhernenkov, H.-Y. N. Holman, S. Lee, L. Chen, E. Rotenberg, T. Weber, Y. L. Goc, M. Boehm, P. Steffens, P. Mutti and J. A. Sethian, Gaussian Processes for Autonomous Data Acquisition at Large-Scale Synchrotron and Neutron Facilities, *Nat. Rev. Phys.*, 2021, 3(10), 685–697, DOI: [10.1038/s42254-021-00345-y](https://doi.org/10.1038/s42254-021-00345-y).
- 30 A. Nicolas, M. Messner and T. Sham, *Initial Development of a High Temperature Life Prediction Method Directly Accounting for Variability in Material Properties*, Argonne National Laboratory, 2020, ANL-ART-203, DOI: [10.2172/1658590](https://doi.org/10.2172/1658590).
- 31 N. Knudde; J. V. Herten; T. Dhaene and I. Couckuyt: A Bayesian Optimization Library Using TensorFlow, *arXiv*, 2017, DOI: [10.48550/arXiv.1711.03845](https://doi.org/10.48550/arXiv.1711.03845).
- 32 Y. Morita, S. Rezaeiravesh, N. Tabatabaei, R. Vinuesa, K. Fukagata and P. Schlatter, Applying Bayesian Optimization with Gaussian Process Regression to Computational Fluid Dynamics Problems, *J. Comput. Phys.*, 2022, 449, 110788, DOI: [10.1016/j.jcp.2021.110788](https://doi.org/10.1016/j.jcp.2021.110788).
- 33 D. De Witte, J. Qing, I. Couckuyt, T. Dhaene, D. Vande Ginste and D. Spina, A Robust Bayesian Optimization Framework for Microwave Circuit Design under Uncertainty, *Electronics*, 2022, 11(14), 2267, DOI: [10.3390/electronics11142267](https://doi.org/10.3390/electronics11142267).
- 34 W. Kobayashi, T. Otsuka, Y. K. Wakabayashi and G. Tei, Physics-Informed Bayesian Optimization Suitable for Extrapolation of Materials Growth, *npj Comput. Mater.*, 2025, 11(1), 36, DOI: [10.1038/s41524-025-01522-8](https://doi.org/10.1038/s41524-025-01522-8).
- 35 D. Khatamsaz, R. Neuberger, A. M. Roy, S. H. Zadeh, R. Otis and R. A. Arróyave, Physics Informed Bayesian Optimization Approach for Material Design: Application to NiTi Shape Memory Alloys, *npj Comput. Mater.*, 2023, 9(1), 221, DOI: [10.1038/s41524-023-01173-7](https://doi.org/10.1038/s41524-023-01173-7).
- 36 V. L. Deringer, A. P. Bartók, N. Bernstein, D. M. Wilkins, M. Ceriotti and G. Csányi, Gaussian Process Regression for Materials and Molecules, *Chem. Rev.*, 2021, 121(16), 10073–10141, DOI: [10.1021/acs.chemrev.1c00022](https://doi.org/10.1021/acs.chemrev.1c00022).
- 37 B. Z. Newby and R. J. Composto, Influence of Lateral Confinement on Phase Separation in Thin Film Polymer Blends, *Macromolecules*, 2000, 33(9), 3274–3282, DOI: [10.1021/ma992092m](https://doi.org/10.1021/ma992092m).
- 38 Studio Encoding Parameters of Digital Television for Standard 4:3 and Wide-Screen16:9 Aspect Ratios, *ITU-R BT.601-7*, International Telecommunication Union, 2011, <https://www.itu.int/rec/R-REC-BT.601-7-201103-I/en>, accessed 2026-03-04.
- 39 M. Noack, J. Sethian, K. Yager, M. Fukuto, P. Zwart, D. Ushizima, H. Krishnan, S. Lee, R. Pandolfi, I. Humphrey, D. Perryman, P. Enfedaque, A. Hexemer, S. Sarker, T. Weber, A. Mehta, R. Li, G. Doerk and M. Boehm, *gpCAM (8.1.9)*, Zenodo, 2024, DOI: [10.5281/zenodo.13993973](https://doi.org/10.5281/zenodo.13993973).
- 40 *scipy.optimize.differential\_evolution—SciPy v1.13.1 Manual*, [https://docs.scipy.org/doc/scipy-1.13.1/reference/generated/scipy.optimize.differential\\_evolution.html](https://docs.scipy.org/doc/scipy-1.13.1/reference/generated/scipy.optimize.differential_evolution.html), accessed 2026-02-06.
- 41 C. Bian, X. Wang, C. Liu, X. Xie and L. Haitao, Impact of Exploration-Exploitation Trade-off on UCB-Based Bayesian Optimization, *IOP Conf. Ser. Mater. Sci. Eng.*, 2021, 1081(1), 012023, DOI: [10.1088/1757-899X/1081/1/012023](https://doi.org/10.1088/1757-899X/1081/1/012023).

