Digital Discovery

Accepted Manuscript

This article can be cited before page numbers have been issued, to do this please use: M. Fan, Y. Wang, G. Vazquez, R. Zhou, I. Karaman, R. Arroyave and X. Qian, *Digital Discovery*, 2026, DOI: 10.1039/D5DD00459D.

ROYAL SOCIETY OF CHEMISTRY

rsc.li/digitaldiscovery

View Article Online
DOI: 10.1039/D5DD00459D

Bayesian Active Learning to Accelerate High Throughput Phase Diagram Exploration

Mingzhou Fan[a,†], Yucheng Wang[a,†], Guillermo Vazquez[b], Ruida Zhou[c], Ibrahim Karaman[b], Raymundo Arróyave[b,d,e,*], Xiaoning Qian[a,f,g,*]

[a] *Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA 77843*

[b] *Department of Materials Science and Engineering, Texas A&M University, College Station, TX, USA 77843*

[c] *Department of Electrical and Computer Engineering, University of California, Los Angeles, CA, USA, 90095*

[d] *J. Mike Walker '66 Department of Mechanical Engineering, Texas A&M University, College Station, TX, USA 77843*

[e] *Wm Michael Barnes '64 Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX 77843*

[f] *Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77843*

[g] *Computing & Data Sciences Directorate, Brookhaven National Laboratory, Upton, NY 11973*

Abstract

Phase diagrams are fundamental for understanding phase stability and guiding the synthesis of new materials. However, constructing high-dimensional phase diagrams through exhaustive CALPHAD (CALculation of PHAse Diagrams) computations remains costly. We introduce a **Bayesian Active Learning for Phase Diagram Discovery (BALPI)** framework that efficiently identifies phase stability regions by adaptively sampling the thermodynamic space using uncertainty-aware acquisition strategies. BALPI integrates Gaussian Process Classifiers and Regressors within two complementary formulations—classification and level-set estimation—and introduces non-myopic Bayesian acquisition functions, including the Soft Mean Objective Cost of Uncertainty (SMOCU) and an extended straddle (e-straddle) criterion. Using CALPHAD-based phase stability predictions as the ground-truth oracle, BALPI achieves accurate reconstruction of phase boundaries with significantly fewer queries than conventional label propagation and label spreading baselines. Results on SiO_2–Al_2O_3–MgO and Ni–Ti–Hf–Cu systems demonstrate that BALPI captures disconnected phase regions and achieves consistent reductions in Bayesian error and computational cost. More importantly, this work establishes BALPI as a general framework for uncertainty-guided phase diagram discovery and highlights the potential of Bayesian active learning to accelerate computational thermodynamics and materials design, through the efficient exploration of the phase stability landscape at much lower costs relative to competing strategies.

Keywords: Bayesian Active Learning, Gaussian Process, Phase Identification, CALPHAD

*Corresponding author emails: raymundo.arroyave@tamu.edu, xqian@ece.tamu.edu
[1] † These authors contributed equally as first authors.

View Article Online
DOI: 10.1039/D5DD00459D

1. Introduction

Phase stability underlies the design and optimization of all materials systems. The equilibrium phase constitution of a system depends on external or internal thermodynamic conditions—temperature, pressure, and composition—and is summarized in *phase diagrams*. From a thermodynamic standpoint, the equilibrium state corresponds to the global constrained minimum of the Gibbs free energy over all phases that may participate in equilibrium. The phase boundaries that comprise the phase diagram separate regions of thermodynamic space associated with distinct equilibrium phase assemblages. These diagrams of stability, in turn, serve as blueprints for understanding microstructural evolution and selecting processing routes that yield desired material properties [1, 2].

The **CALPHAD** (CALculation of PHAse Diagrams) framework has become the *de facto* standard for constructing phase diagrams of complex alloy systems [3, 4]. By parameterizing the Gibbs free energies of all relevant phases and minimizing the total free energy under specified conditions, CALPHAD enables the rigorous prediction of equilibrium phase stability across a desired thermodynamic space. These free energy functions are constructed by combining experimental phase stability measurements, thermo-chemical data, and first-principles calculations [5]. While powerful, this process remains computationally and experimentally intensive—particularly for multi-component systems where exhaustive sampling of the compositional and thermal space is infeasible. Moreover, traditional grid-based exploration strategies fail to leverage the fact that certain compositions are disproportionately informative for thermodynamic model calibration and phase boundary discovery.

From grid-based sampling to active learning. Active Learning (AL) offers a principled approach to efficiently explore design spaces by sequentially selecting the most informative data points to label [6]. In materials discovery, AL has proven effective in accelerating experimental design and phase diagram mapping [7–11]. For instance, the Phase-Mapper project employed AL to enhance phase identification in combinatorial thin-film libraries through adaptive sampling guided by X-ray diffraction feedback [12, 13]. More recent studies applied label propagation and label spreading models to automate phase-diagram reconstruction [11]. However, these methods rely on heuristic uncertainty measures or KNN-based surrogates whose performance is highly sensitive to hyperparameter choices. In contrast, *Gaussian Processes* (GPs) provide a Bayesian, non-parametric alternative that enables closed-form uncertainty quantification and empirical Bayes hyperparameter optimization [14, 15]. Building on this foundation, recent work has demonstrated how embedding physics-informed priors into the mean function of GP classifiers can improve constraint-aware sampling and accelerate phase stability discovery under limited data conditions [16].

Bayesian Active Learning for CALPHAD-guided phase discovery. Here, we introduce a unified framework—**Bayesian Active Learning for Phase Identification (BALPI)**—that couples uncertainty-aware sampling with CALPHAD-based phase computations. Rather than re-parameterizing thermodynamic models, BALPI treats CALPHAD (implemented via Thermo-Calc) as a ground-truth oracle that returns either discrete phase labels or continuous phase-fraction scores. The learner iteratively queries new compositions in thermodynamic space to minimize global predictive uncertainty in the phase stability landscape. The use of a CALPHAD oracle mimics an experimental workflow in which phase stability is assessed composition-by-composition via synthesis, followed by characterization of the resulting phase constitution using microscopy, spectroscopy, or diffraction techniques.

We formulate phase discovery under two complementary views:

1. **Classification:** learning the stability domains of distinct phases as a multi-class decision problem using a Gaussian Process Classifier (GPC);
2. **Level-set estimation:** modeling a continuous latent function corresponding to phase-fraction or stability scores using a Gaussian Process Regressor (GPR).

Unlike prior active learning approaches for phase diagram construction that rely primarily on discrete phase labels and myopic uncertainty sampling [7–11, 17], BALPI unifies classification and level-set estimation formulations within

View Article Online
DOI: 10.1039/D5DD00459D

a single Bayesian perspective. In particular, BALPI supports queries of either binary phase indicators or continuous phase-fraction or stability scores, enabling the use of both GPC and GPR models coupled with non-myopic acquisition functions such as Soft Mean Objective Cost of Uncertainty (SMOCU) [18, 19] and extended straddle. Importantly, BALPI is agnostic to the source of phase information: while CALPHAD equilibrium computations are used in this work as a reproducible oracle for benchmarking, the framework is directly applicable to experimental workflows in which phase stability is determined via synthesis and characterization.

For adaptive sampling, we evaluate a suite of Bayesian acquisition functions, including: (i) Maximum Entropy Search (MES), which selects points with maximal expected information gain [20]; (ii) Bayesian Active Learning by Disagreement (BALD), which targets epistemic uncertainty via disagreement in posterior predictive distributions [21]; (iii) a soft Mean Objective Cost of Uncertainty (SMOCU) criterion for non-myopic, outcome-aware exploration [18, 19]; and (iv) an extended straddle (e-straddle) function that balances phase-boundary refinement with global uncertainty reduction [22].

Contributions and impact. Our contributions are threefold:

- We propose a modular Bayesian active learning platform (BALPI) for uncertainty-aware exploration of CALPHAD-computed phase diagrams.
- We formulate phase discovery as either a classification or a level-set estimation problem, combining GPC/GPR models with principled Bayesian acquisition functions.
- We empirically demonstrate that BALPI outperforms conventional label propagation and label spreading baselines on representative ternary and quaternary systems, achieving faster convergence and resolving disconnected phase regions that are often missed by deterministic methods.

Importantly, and beyond CALPHAD-related applications, BALPI establishes a general framework for autonomous phase diagram discovery where querying phase information—experimentally or via simulation—is costly. The combination of Gaussian process surrogates and Bayesian acquisition functions offers a robust pathway toward interpretable, uncertainty-aware materials exploration.

2. Related Work

Active learning (AL) has emerged as a powerful strategy for accelerating the construction of phase diagrams and phase mapping in both computational and experimental materials science. Early efforts such as the *Phase-Mapper Project* [12, 13] demonstrated how machine learning can guide experimental synthesis and phase identification from high-throughput X-ray diffraction data. Building upon these ideas, Terayama *et al.* [11] introduced an uncertainty-sampling strategy for efficient phase diagram construction by integrating label propagation and label spreading models with CALPHAD-based thermodynamic calculations.

More recent studies have incorporated explicit Bayesian or thermodynamics-informed models to further improve sampling efficiency. Ament *et al.* [7] proposed a hierarchical Bayesian active learning framework capable of autonomously discovering nonequilibrium phase diagrams from experimental data. Abranches *et al.* [9] developed a thermodynamics-informed Gaussian process (GP) active learning approach to estimate activity coefficients and build phase diagrams, emphasizing the integration of domain-specific priors. Dai *et al.* [8] formulated phase diagram exploration as an active learning problem with adaptive sampling driven by model uncertainty, while Deffrennes *et al.* [10] combined supervised machine learning with CALPHAD computations to directly predict phase boundaries. In parallel, Lookman *et al.* [17] demonstrated that Gaussian process regression and Bayesian inference provide a flexible and interpretable foundation for modeling phase diagrams and efficiently sampling phase boundaries.

View Article Online
DOI: 10.1039/D5DD00459D

Table 1: Summary of notation used in the Bayesian Active Learning for Phase Diagram Discovery (BALPI) framework.

Symbol	Description
$\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$	Composition vector in d-dimensional design space
T, P	Temperature and pressure defining thermodynamic conditions
$y \in \{0, 1\}$	Binary phase indicator (1: phase of interest, 0: otherwise)
$s \in [0, 1]$	Continuous stability score or phase fraction (e.g., CALPHAD output)
$f(\mathbf{x})$	Latent phase-stability function modeled by the GP surrogate
$\mathcal{D}_t$	Dataset at iteration t: $\{(\mathbf{x}_i, y_i)\}_{i=1}^{N_t}$
$K(\mathbf{x}, \mathbf{x}')$	Covariance (kernel) function defining GP prior correlation
a, l	Kernel hyperparameters: amplitude and length scale
$\mu(\mathbf{x}), \sigma^2(\mathbf{x})$	GP posterior mean and variance at $\mathbf{x}$
$U(\mathbf{x})$	Bayesian acquisition (utility) function guiding query selection
$U_{\text{SMOCU}}(\mathbf{x})$	Soft Mean Objective Cost of Uncertainty (decision-theoretic utility)
$U_{\text{e-straddle}}(\mathbf{x})$	Extended straddle utility (heuristic boundary-based utility)
t	Active learning iteration.
$g(\cdot)$	Transformation in e-straddle reward term ($g(r) = r$, r^2, or $\sqrt{r}$)
β	Hyperparameter balancing exploration and exploitation (e-straddle)
k	Softness parameter in S-MOCU controlling LogSumExp smoothness
$\mathcal{D}_T$	Final dataset after T active learning iterations
$p(y \mid \mathbf{x}, \mathcal{D})$	Predictive phase probability from GP surrogate
$\text{Var}(f(\mathbf{x}) \mid \mathcal{D})$	Posterior variance of latent function (epistemic uncertainty)
$\mathbf{x}_{t+1}$	Next query selected by maximizing $U(\mathbf{x})$
τ	Decision threshold for phase boundary (typically $t = 0.5$)

Despite these advances, most existing frameworks rely on myopic acquisition strategies or a single modeling formulation (typically classification), limiting their ability to identify disconnected or multimodal phase-stability regions. Our proposed **Bayesian Active Learning for Phase Identification (BALPI)** framework extends this line of research by (i) unifying classification and level-set estimation formulations under a single Bayesian perspective, (ii) introducing non-myopic and differentiable acquisition functions such as the *Soft Mean Objective Cost of Uncertainty (SMOCU)* and *extended straddle (e-straddle)*, and (iii) demonstrating superior convergence and robustness on both glass-ceramic and shape-memory alloy systems. BALPI thus generalizes prior AL paradigms into a modular, uncertainty-aware framework for the autonomous exploration of CALPHAD-derived thermodynamic spaces.

View Article Online
DOI: 10.1039/D5DD00459D

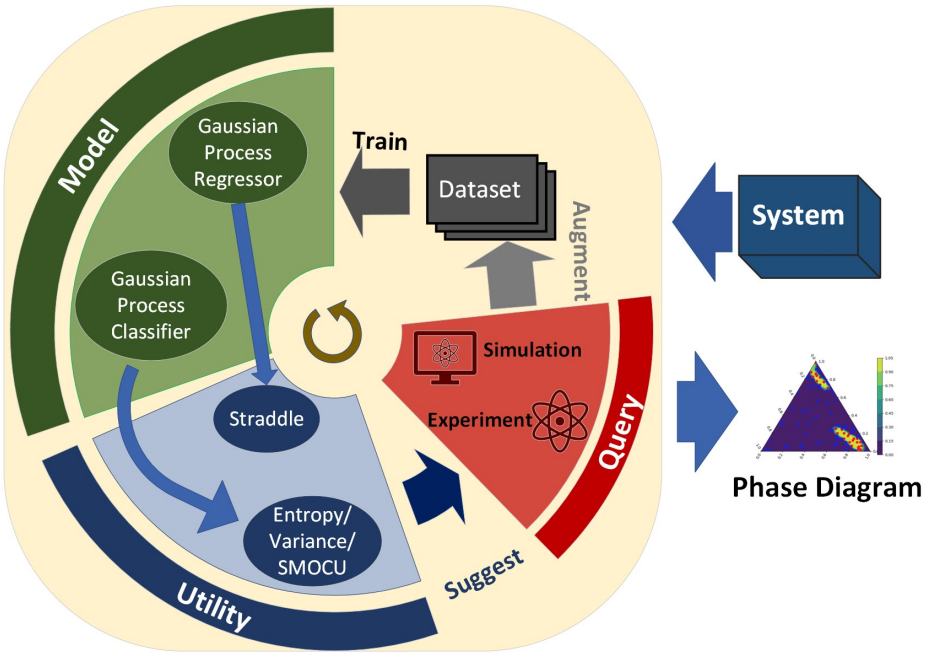

Figure 1: Schematic of the **BALPI** (Bayesian Active Learning for Phase Diagram Discovery) workflow. The framework integrates *Model*, *Utility*, and *Query* modules: a Gaussian Process Classifier or Regressor serves as the surrogate model; Bayesian utilities such as Straddle, Entropy, Variance, or SMOCU guide sampling; and CALPHAD simulations or experiments provide new phase observations to augment the dataset, iteratively refining the phase diagram.

3. Methods

The proposed **Bayesian Active Learning for Phase Diagram Discovery (BALPI)** framework provides a unified workflow for automated phase diagram exploration by combining CALPHAD-based thermodynamic modeling with Gaussian Process (GP)–based surrogate learning and Bayesian decision-making. As illustrated in Figure 1, BALPI operates as a closed-loop system consisting of three modules: *Model*, *Utility*, and *Query*. In each iteration, the GP surrogate (either a classifier or a regressor) is trained on existing CALPHAD data—serving as a proxy for a fully experimental or computational campaign. The Bayesian acquisition function is then used to compute the expected information gain across all candidate compositions, and the most informative sample is selected for evaluation—either by querying the CALPHAD oracle or by performing a phase-stability experiment. The resulting phase label or stability score is appended to the dataset, and the loop continues until convergence in uncertainty or exhaustion of the sampling budget. This design allows BALPI to systematically map phase stability regions with minimal computational cost, serving as a general-purpose platform for uncertainty-guided exploration of thermodynamic systems.

From a technical perspective, the BALPI workflow can be interpreted as a Bayesian decision-making process built on Gaussian Process (GP) surrogates. Each module in Figure 1 corresponds to a distinct mathematical component: the *Model* stage defines a GP prior over the latent thermodynamic function $f(\mathbf{x})$, the *Utility* stage evaluates an acquisition function $U(\mathbf{x})$ that quantifies the expected benefit — such as information gain, uncertainty reduction, or decision-quality improvement — from evaluating a new query, and the *Query* stage provides ground-truth phase data from CALPHAD or experiments to update the dataset. This modular decomposition allows the active learning process to be formalized within a unified probabilistic framework, where the objective is to minimize the epistemic uncertainty of the surrogate model while efficiently identifying phase boundaries. The following subsections provide the mathematical formulation of this framework, including (i) the problem setup, (ii) the Gaussian Process surrogate models for phase prediction, (iii) the Bayesian acquisition functions used to guide exploration, and (iv) the complete

View Article Online
DOI: 10.1039/D5DD00459D

active learning loop that defines the BALPI algorithm.

3.1. Problem formulations

The goal of BALPI is to efficiently discover the phase stability landscape of a multi-component material system within a specified thermodynamic space. Let the system be parameterized by composition $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$, temperature T, and pressure P. Each query to the thermodynamic model or experiment returns a phase indicator, which can be represented as either a discrete label

$$y \in \{0, 1\},$$

where $y = 1$ indicates that the queried design $\mathbf{x}$ belongs to the phase of interest and $y = 0$ otherwise, or a continuous score

$$s \in [0, 1],$$

such as the predicted phase fraction or stability probability obtained from CALPHAD-based simulations (e.g., Thermo-Calc). The dataset at iteration t is thus denoted as

$$\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_t} \quad \text{or} \quad \mathcal{D}_t = \{(\mathbf{x}_i, s_i)\}_{i=1}^{N_t},$$

depending on the availability of discrete or continuous outputs.

The objective of BALPI is to construct a surrogate model that accurately approximates the phase stability function across $\mathcal{X}$ while minimizing the number of expensive CALPHAD evaluations or experimental measurements. This objective can be formalized as minimizing the total model uncertainty over the domain:

$$\min_{\mathcal{D}_T} \mathbb{E}_{\mathbf{x}\in\mathcal{X}}[\mathrm{Var}(f(\mathbf{x}) \mid \mathcal{D}_T)],$$

where $f(\mathbf{x})$ denotes the latent function representing the thermodynamic state (e.g., the phase indicator score).

At each iteration, BALPI actively selects the next query $\mathbf{x}_{t+1}$ according to a Bayesian acquisition function $U(\mathbf{x})$:

$$\mathbf{x}_{t+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} U(\mathbf{x}),$$

where $U(\mathbf{x})$ quantifies the expected information gain or uncertainty reduction that would result from evaluating the phase at $\mathbf{x}$. This active querying strategy enables targeted sampling near phase boundaries, where uncertainty in phase stability is typically the highest, thereby accelerating convergence of the phase diagram.

In this study, we focus on the ternary system SiO_2–Al_2O_3–MgO as a representative case for demonstrating the efficiency of the proposed framework. However, the formulation is general and applicable to arbitrary multicomponent systems and thermodynamic conditions.

Feature preprocessing and compositional constraints. In all experiments, the model inputs are composition fractions, so the feasible design space is a compositional simplex rather than an unconstrained hypercube. We therefore parameterize the surrogate directly in terms of independent composition coordinates and enforce the sum-to-one constraint analytically. For the ternary SiO_2–Al_2O_3–MgO case for example, the independent inputs include (x, y) and the third coordinate is recovered as $z = 1 - x - y$, with only feasible points satisfying $x, y, z \in [0, 1]$ retained. For the quaternary NiTiHfCu case, compositions are parameterized as $Cu_xHf_yNi_{z/2}Ti_{z/2}$ with $x + y + z = 1$, so the independent variables again lie on a simplex that uniquely determines the full composition. Because these variables are already expressed as normalized mole fractions on [0, 1], we do not introduce an additional z-score or min-max feature scaling step beyond enforcing the compositional feasibility constraint.

View Article Online
DOI: 10.1039/D5DD00459D

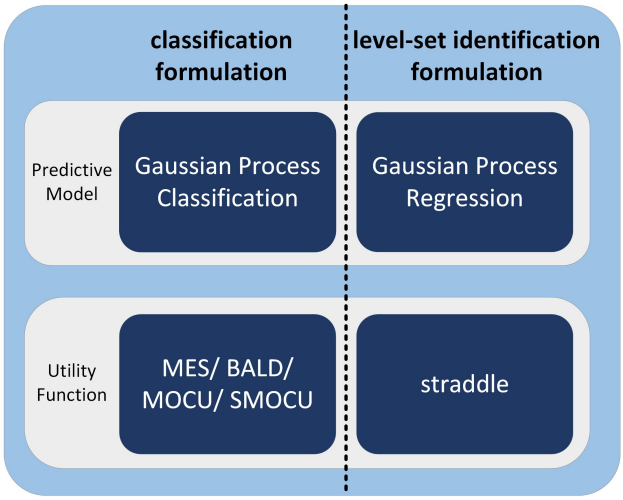

Figure 2: Predictive models and Bayesian utility functions used in the two phase identification formulations of the BALPI framework. The classification formulation employs Gaussian Process Classification with MES, BALD, MOCU, or SMOCU utilities, while the level-set formulation uses Gaussian Process Regression with the straddle utility.

Due to differences in problem formulation, we employ Gaussian Process Classification (GPC) and Gaussian Process Regression (GPR) as predictive models for the classification and regression-based versions of BALPI, respectively. Each formulation is paired with its own set of Bayesian utility functions. The predictive models and associated utility functions are summarized in Figure 2, and the full BALPI workflow is illustrated in Figure 1. We describe both components in more detail below.

3.2. *Gaussian Process Surrogates*

In Bayesian Active Learning, the predictive surrogate plays a central role by providing both mean estimates and principled uncertainty quantification required to compute the acquisition utility. Gaussian Processes (GPs) are a natural choice for this purpose due to their analytical tractability, interpretability, and inherent ability to model predictive uncertainty. They have been widely adopted as surrogate models in Bayesian optimization and Bayesian experimental design [23], particularly in problems with moderate input dimensionality such as the phase identification tasks considered in this work. We employ two GP variants: the Gaussian Process Classifier (GPC) for phase classification and the Gaussian Process Regressor (GPR) for level-set estimation. Both models share a common prior formulation, but differ in their likelihood models.

Gaussian Process Classifier (GPC). For classification-based phase identification, we define a data set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$ with labels $y_i \in \{0, 1\}$ that indicate the presence or absence of a specific phase. The GPC assumes a latent function $f(\mathbf{x})$ drawn from a Gaussian process prior:

$$p(f) = \mathcal{GP}(0, K(\mathbf{x}, \mathbf{x}')), \tag{1}$$

where $K(\cdot, \cdot)$ is a positive-definite kernel function encoding the smoothness of f. Each label y_i follows a Bernoulli likelihood with a sigmoid link function:

$$p(y_i \mid f(\mathbf{x}_i)) = \sigma(f(\mathbf{x}_i))^{y_i}[1 - \sigma(f(\mathbf{x}_i))]^{1-y_i}, \tag{2}$$

where $\sigma(\cdot)$ denotes the logistic sigmoid. Because this likelihood is non-Gaussian, the posterior $p(f \mid \mathcal{D})$ is analytically intractable. Approximate inference methods such as Laplace Approximation (LA) [24] and Expectation Propagation (EP) [25] are therefore employed. We adopt EP in this work for its superior accuracy in approximating the latent

View Article Online
DOI: 10.1039/D5DD00459D

function posterior through iterative local Gaussian refinement. The resulting predictive distribution $p(y_* \mid \mathbf{x}_*, \mathcal{D})$ provides class probabilities and corresponding predictive uncertainty, which serve as the foundation for evaluating uncertainty-aware acquisition functions such as SMOCU.

Gaussian Process Regressor (GPR). For level-set estimation and continuous phase-score modeling, we employ a Gaussian Process Regressor defined on the dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i^c)\}_{i=1}^N$, where each observation follows

$$y_i^c = f(\mathbf{x}_i) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_n^2),$$

with σ_n^2 representing Gaussian observation noise. Assuming the same prior as in Eq. (1), the posterior predictive distribution at any new input $\mathbf{x}_*$ is Gaussian:

$$p(f_* \mid \mathbf{x}_*, \mathbf{X}, \mathbf{y}) = \mathcal{N}(\mu(\mathbf{x}_*), \sigma^2(\mathbf{x}_*)), \tag{3}$$

with closed-form expressions for the mean and variance:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \tag{4}$$

$$\sigma^2(\mathbf{x}_*) = K(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*, \tag{5}$$

where $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$, $\mathbf{y} = [y_1^c, \ldots, y_N^c]^\top$, and the kernel matrix $\mathbf{K} = [\mathbf{k}_1, \ldots, \mathbf{k}_N]$, $\mathbf{k}_j = [K(\mathbf{x}_1, \mathbf{x}_j), \ldots, K(\mathbf{x}_N, \mathbf{x}_j)]^\top$, $\mathbf{k}_* = [K(\mathbf{x}_1, \mathbf{x}_*), \ldots, K(\mathbf{x}_N, \mathbf{x}_*)]^\top$. The predictive variance $\sigma^2(\mathbf{x}_*)$ naturally captures the epistemic uncertainty of the model and provides the quantitative basis for active sampling in level-set estimation tasks.

Kernel Function and Hyperparameter Estimation. The covariance kernel $K(\mathbf{x}, \mathbf{x}')$ specifies the correlation structure of the GP prior and governs the smoothness of the learned phase landscape. Among standard choices such as the Radial Basis Function (RBF), Matérn, and linear kernels, we employ the RBF kernel for its interpretability and smooth generalization:

$$K_{a,l}(\mathbf{x}, \mathbf{x}') = a \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{l}\right), \tag{6}$$

where a and l denote the kernel amplitude and length-scale hyperparameters, respectively. These hyperparameters are estimated using empirical Bayes by maximizing the marginal likelihood:

$$p(\mathbf{y} \mid \mathbf{X}) = \frac{1}{(2\pi)^{\frac{N}{2}} |\mathbf{K}_{a,l}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} \mathbf{y}^\top \mathbf{K}_{a,l}^{-1} \mathbf{y}\right), \tag{7}$$

where $\mathbf{K}_{a,l}$ is the kernel matrix constructed from $K_{a,l}(\mathbf{x}, \mathbf{x}')$. The resulting GP surrogate thus provides both mean predictions and calibrated uncertainty estimates, forming the probabilistic foundation for the acquisition functions described in the following section.

3.3. Bayesian Acquisition Functions

The acquisition function determines the sampling policy in Bayesian Active Learning (BAL) by quantifying the expected information gain or uncertainty reduction from evaluating a new point. In this work, we focus on two complementary acquisition strategies: a decision-theoretic utility based on the *Soft Mean Objective Cost of Uncertainty* (S-MOCU) and a heuristic, uncertainty-driven criterion termed *extended straddle* (e-straddle). Both approaches rely on the predictive uncertainty provided by the Gaussian Process surrogates described in Section 3.2.

View Article Online
DOI: 10.1039/D5DD00459D

Soft Mean Objective Cost of Uncertainty (S-MOCU). The Mean Objective Cost of Uncertainty (MOCU) framework [18, 19] quantifies the expected degradation in performance caused by model uncertainty. In the context of phase identification, MOCU measures the expected loss in decision quality that arises from uncertainty in the phase stability classifier across the design space. To minimize this loss, the MOCU-based acquisition policy selects the next query that maximally reduces global uncertainty in a one-step look-ahead manner.

We adopt a **Soft MOCU (S-MOCU)** formulation that relaxes the original non-differentiable MOCU objective by replacing the max operator with a smooth `LogSumExp` approximation. This results in a strictly concave and differentiable utility function, enabling efficient gradient-based optimization over the design space.

The S-MOCU utility is defined as

$$U_{\mathrm{SMOCU}}(\mathbf{x}) = \mathbb{E}_{\mathbf{x}'}\Big[\mathbb{E}_{y|\mathbf{x}}\Big[\tfrac{1}{k}\,\mathrm{LogSumExp}(k\,p(y \mid \mathbf{x}', \mathcal{D}, \mathbf{x}, y)) - \tfrac{1}{k}\,\mathrm{LogSumExp}(k\,p(y \mid \mathbf{x}', \mathcal{D}))\Big]\Big]. \tag{8}$$

where $\mathbf{x}$ is the candidate composition to query, $\mathbf{x}'$ spans all points in the design space, and $p(y|\mathbf{x}', \mathcal{D}, \mathbf{x}, y)$ denotes the predictive distribution after a hypothetical observation y at $\mathbf{x}$. The parameter $k > 0$ controls the smoothness of the LogSumExp relaxation: as $k \to \infty$, the expression converges to the original MOCU objective, while smaller k yields smoother gradients. By softening the discontinuous maximum, S-MOCU preserves the decision-theoretic interpretation of MOCU while offering improved numerical stability and more efficient optimization. Its one-step look-ahead nature ensures globally informed exploration that adaptively focuses sampling in regions of high phase uncertainty.

Extended Straddle (e-straddle). While S-MOCU provides a principled decision-theoretic approach, the straddle criterion [26] offers a computationally efficient, heuristic alternative particularly suited for level-set estimation tasks. The classical *straddle* utility,

$$U_{\mathrm{straddle}}(\mathbf{x}) = -|\mu(\mathbf{x})| + 1.96\,\sigma(\mathbf{x}), \tag{9}$$

can be interpreted as an upper-confidence-bound (UCB) rule that balances proximity to the decision boundary (exploitation) and prediction uncertainty (exploration).

We generalize this formulation and propose the **extended straddle (e-straddle)** utility:

$$U_{\text{e-straddle}}(\mathbf{x}) = -\,g(|\mu(\mathbf{x}) - \tau|) + \beta\,\sigma(\mathbf{x}), \tag{10}$$

where τ denotes the decision threshold (typically $\tau = 0.5$), β is a tunable coefficient controlling the exploration–exploitation balance, and $g(\cdot)$ is a transformation function applied to the reward term. The transformation modulates the sampling preference with respect to the distance from the phase boundary. Specifically:

- $g(r) = r$ recovers the standard straddle criterion;
- $g(r) = r^2$ increases the penalty for points farther from the boundary, encouraging broader exploration;
- $g(r) = \sqrt{r}$ reduces this penalty, refining sampling near the boundary.

By adjusting $g(\cdot)$, the e-straddle framework flexibly interpolates between exploratory and exploitative behaviors, allowing efficient and interpretable control of the active sampling dynamics. While heuristic in nature, e-straddle remains computationally lightweight and performs competitively against more complex Bayesian criteria such as S-MOCU.

View Article Online
DOI: 10.1039/D5DD00459D

Comparison and Complementarity. The two acquisition functions are conceptually complementary: S-MOCU explicitly minimizes the expected global cost of uncertainty and thus offers a theoretically grounded approach to phase-space exploration, whereas e-straddle provides a simpler yet effective heuristic for rapidly refining phase boundaries. Together, these two strategies enable both globally aware and locally adaptive sampling within the proposed BAL framework.

3.4. Active Learning Workflow Implementation

The proposed Bayesian Active Learning for Phase Identification (BALPI) framework integrates Gaussian Process surrogates with Bayesian acquisition strategies to adaptively construct phase diagrams with minimal thermodynamic evaluations. Figure 1 provides an overview of the iterative workflow, and Algorithm 1 summarizes the full procedure.

Workflow Overview. The BALPI workflow consists of four main stages:

1. **Initialization.** A small set of initial sampling points $\{\mathbf{x}_i\}_{i=1}^{N_0}$ is selected using a low-discrepancy design (e.g., Latin hypercube sampling) or uniformly over the design space $\mathcal{X}$. Phase labels or scores $\{y_i\}$ are obtained via CALPHAD calculations or experiments, forming the initial dataset $\mathcal{D}_0$.
2. **Surrogate Model Training.** A Gaussian Process surrogate (GPC or GPR, depending on whether the task involves classification or level-set estimation) is trained on $\mathcal{D}_t$. The surrogate provides predictive mean $\mu(\mathbf{x})$ and uncertainty $\sigma^2(\mathbf{x})$ for all candidate compositions.
3. **Acquisition and Query Selection.** The acquisition function $U(\mathbf{x})$ is evaluated across the design space using either the S-MOCU utility for global uncertainty reduction or the e-straddle utility for efficient boundary refinement. The next query point is chosen as

$$\mathbf{x}_{t+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} U(\mathbf{x}).$$

   The new phase observation y_{t+1} is then obtained and appended to the dataset: $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_{t+1}, y_{t+1})\}$.
4. **Update and Convergence.** The surrogate is retrained with the updated dataset $\mathcal{D}_{t+1}$, and the process repeats until a predefined stopping criterion is met, such as (i) the maximum predictive uncertainty falls below a tolerance, or (ii) the number of allowed evaluations T is reached.

This iterative loop efficiently allocates computational or experimental resources to the most informative regions of the thermodynamic space—typically near multi-phase boundaries—resulting in a rapidly convergent and uncertainty-aware phase mapping.

Algorithm 1 Bayesian Active Learning for Phase Identification (BALPI)

Require: Design space $\mathcal{X}$, initial dataset $\mathcal{D}_0$, surrogate model (GPC or GPR), acquisition function $U(\mathbf{x})$, maximum iterations T

for $t = 0$ to $T - 1$ **do**
  Train GP surrogate on $\mathcal{D}_t$
  Evaluate acquisition function $U(\mathbf{x})$ across $\mathcal{X}$
  Select next query: $\mathbf{x}_{t+1} = \arg\max_{\mathbf{x}\in\mathcal{X}} U(\mathbf{x})$
  Obtain observation y_{t+1} from CALPHAD or experiment
  Update dataset: $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_{t+1}, y_{t+1})\}$
end for
return Final surrogate model and predicted phase map

View Article Online
DOI: 10.1039/D5DD00459D

Implementation and Integration with CALPHAD. In this work, the BALPI framework interfaces directly with the Thermo-Calc software package to obtain CALPHAD-based phase stability data. The surrogate and acquisition functions are implemented in Python using `NumPy`, allowing efficient training, batched evaluation, and uncertainty propagation. The modular design of the platform enables future extensions to higher-dimensional or multi-fidelity settings, as well as integration with experimental automation systems.

4. Results and Discussion

We evaluate the proposed Bayesian Active Learning for Phase Identification (BALPI) framework on two representative material systems: (1) the SiO_2–Al_2O_3–MgO glass-ceramic glazes system, and (2) the NiTiHfCu shape memory alloy system. These systems highlight different aspects of the framework: the former involves the discovery of well-separated phase stability regions (Spinel and Mullite phases) within a ternary diagram, whereas the latter features multiple disconnected stability domains for the BCC–B2 phase within a more complex composition space. Across both studies, we compare BALPI with prior non-Bayesian approaches based on Label Propagation (LP) and Label Spreading (LS) [11], as well as with other Bayesian active learning utilities including Maximum Entropy Search (MES) [20], Bayesian Active Learning by Disagreement (BALD) [21], and Mean Objective Cost of Uncertainty (MOCU) [18, 19]. Performance is evaluated using the *Bayesian error* metric, which measures the expected phase identification uncertainty over the design space.

4.1. Evaluation Metric

To compare methods consistently, we report the *Bayesian error*, which measures the expected misclassification probability under the surrogate's predictive distribution:

$$\begin{aligned}\text{BayesErr} = \mathbb{E}_{\mathbf{x}\in\mathcal{X}}\Big[&p_{\mathcal{D}}(y=0\mid\mathbf{x})\,\mathbb{I}(y_{\text{true}}=1)\\ &+\ p_{\mathcal{D}}(y=1\mid\mathbf{x})\,\mathbb{I}(y_{\text{true}}=0)\Big].\end{aligned} \tag{11}$$

In practice, we evaluate (11) as an empirical average over a dense composition grid.

Level-set evaluation. For the level-set formulation with GPR, the surrogate models a latent continuous field $f(\mathbf{x}) \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$ and a phase boundary threshold τ. We convert the regression output to phase probabilities via

$$p_{\mathcal{D}}(y=1\mid\mathbf{x}) = \mathbb{P}(f(\mathbf{x}) \geq \tau), \tag{12}$$

$$p_{\mathcal{D}}(y=0\mid\mathbf{x}) = 1 - p_{\mathcal{D}}(y=1\mid\mathbf{x}), \tag{13}$$

These probabilities are then used in (11) for a fair comparison with classification-based methods.

4.2. Phase Identification for the SiO_2–Al_2O_3–MgO Glass-Ceramic System

We first test the performance of BALPI on the SiO_2–Al_2O_3–MgO ternary system, where the objective is to identify the regions of the Spinel and Mullite phases. Both phases are of technological importance due to their high-temperature stability and chemical resistance [27, 28]. The composition space is parameterized as $(SiO_2)_x$–$(Al_2O_3)_y$–$(MgO)_z$ with the compositional constraint $x + y + z = 1$. Each query returns a binary phase label, where $y = 1$ corresponds to the phase of interest (Spinel or Mullite) and $y = 0$ otherwise. The ground-truth phase diagram used for validation is shown in Fig. 3a.

View Article Online
DOI: 10.1039/D5DD00459D

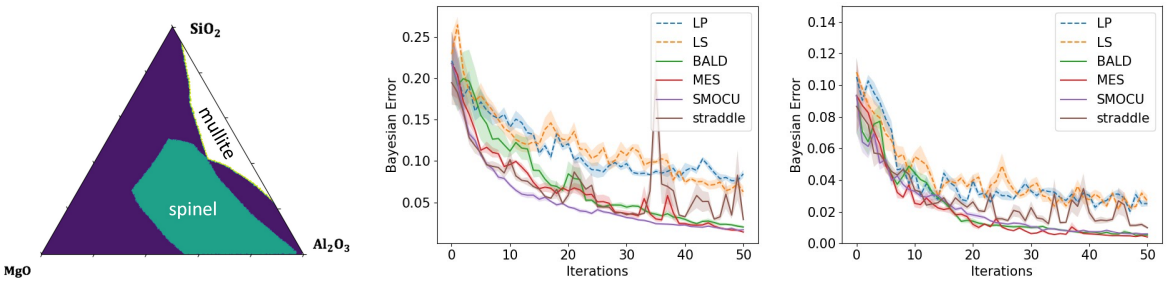

Target phase diagram. Spinel identification. Mullite identification.

Figure 3: (a) Ground-truth phase diagram of $(SiO_2)_x$–$(Al_2O_3)_y$–$(MgO)_z$ system [28]. (b–c) Bayesian error over active learning iterations for Spinel and Mullite phase identification. The proposed BALPI methods achieve significantly lower Bayesian errors than the benchmark LP and LS approaches.

Figures 3b and 3c report the Bayesian error of the competing methods for Spinel and Mullite identification, respectively. In both cases, the proposed BALPI methods achieve substantially lower Bayesian error than LP and LS throughout the active learning iterations. Among all acquisition strategies, the SMOCU-guided BAL exhibits the fastest error decay, consistently outperforming MES and other utilities. This advantage stems from SMOCU's one-step-look-ahead formulation, which explicitly accounts for the global uncertainty reduction effect of each query. Although the *e-straddle* utilities do not always achieve the best performance, they maintain a steady improvement trend and perform markedly better than the non-Bayesian LP and LS baselines. These results demonstrate that the BALPI framework effectively balances exploration and exploitation, even in relatively simple ternary systems.

4.3. BCC–B2 Phase Identification for the NiTiHfCu *Shape Memory Alloy System*

We next evaluate BALPI on the quaternary NiTiHfCu shape memory alloy (SMA) system, a material of increasing interest for high temperature actuation applications [29, 30]. Recent work has shown that NiTiHfCu alloys offer a promising route to achieve high transformation temperatures ($T_{\text{trans}} > 250$ °C) while reducing the large thermal hysteresis that historically has limited the cyclic stability and actuator efficiency of conventional NiTiHf systems [31]. In particular, Cu additions of 5–15 at.% have been shown to promote the formation of B19 orthorhombic martensite in place of B19', enhancing crystallographic compatibility and dramatically reducing hysteresis. However, this benefit is counterbalanced by the strong sensitivity of transformation temperatures to both Cu and Ni content, as well as the potential for brittle precipitate formation in Hf-rich regimes. As demonstrated in our recent constrained Bayesian optimization campaign [31], the thermodynamic design space of NiTiHfCu is extremely sparse, with only a narrow and disconnected manifold yielding compositions that simultaneously meet target martensite start temperatures (250–350 °C) and hysteresis constraints (¡50 °C). The BCC–B2 austenite phase is of particular importance, as it is the only parent phase from which a reversible martensitic transformation can occur. Accordingly, the goal is to localize and characterize regions of high BCC–B2 phase fraction at 800 K with minimal queries to CALPHAD (or, more critically, experimental) oracles. This presents a stringent test for BALPI, which must balance global exploration and local refinement to recover small, topologically fragmented stability domains embedded in a high-dimensional composition-processing space.

The composition space in this system is parameterized as $Cu_xHf_yNi_{z/2}Ti_{z/2}$, subject to $x+y+z=1$. As mentioned above, this system is particularly challenging because the BCC–B2 phase domains are disjoint and separated by multi-phase regions [32–34]. In this work, we perform two sets of experiments: (i) a **binary phase classification**, where compositions with BCC–B2 fraction > 0.8 are labeled as $y = 1$, and (ii) a **continuous level-set estimation**, where the phase stability score itself is directly queried. The corresponding ground-truth phase maps are shown in Figs. 4a and 4b.

View Article Online
DOI: 10.1039/D5DD00459D

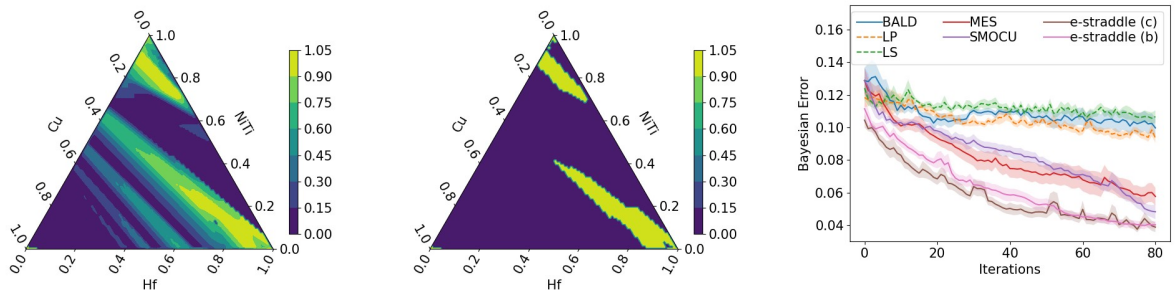

(a) Continuous phase fraction map. (b) Binary classification map. (c) Bayesian error vs. iteration.

Figure 4: (a–b) Target BCC–B2 phase maps for the NiTiHfCu system at 800 K. (c) Bayesian error over active learning iterations showing that all BALPI-based methods outperform LP and LS baselines. The continuous-query *e-straddle(c)* achieves the fastest convergence.

	e-straddle(b)			*e-straddle(c)*		
Transformation	$g(r) = r$	$g(r) = r^2$	$g(r) = \sqrt{r}$	$g(r) = r$	$g(r) = r^2$	$g(r) = \sqrt{r}$
$\beta = 0.5$	0.0492±0.0218	0.0493±0.0143	0.0794±0.0419	0.0479±0.0224	0.0442±0.0106	0.0510±0.0157
$\beta = 1$	**0.0403**±0.0067	0.0542±0.0218	0.0660±0.0408	0.0452±0.0134	0.0526±0.0167	0.0396±0.0134
$\beta = 2$	0.0482±0.0100	0.0524±0.0073	0.0423±0.0096	0.0561±0.0137	0.0515±0.0134	**0.0388**±0.0094

Table 2: Final Bayesian error of the *e-straddle* utility with different transformation functions $g(\cdot)$ and hyperparameters β. Bold values denote the best performance for binary and continuous query functions, respectively.

As illustrated in Fig. 4c, all methods under the BALPI framework achieve lower Bayesian errors than LS and LP for the same query budget. Notably, the *e-straddle* utilities outperform SMOCU and MES in this task, especially in the early iterations where phase boundary localization is critical. Among the two e-straddle formulations, the continuous-query version (*e-straddle(c)*) achieves faster convergence than its binary counterpart (*e-straddle(b)*).

Figure 5 visualizes the queried compositions after 80 iterations for representative methods. Benchmark approaches such as label propagation (LP) tend to oversample large regions of low uncertainty, often missing smaller—and practically critical—pockets of BCC–B2 stability. In contrast, BALPI variants guided by SMOCU and e-straddle successfully identify multiple disconnected phase regions within the sparse stability landscape. This improved sampling efficiency arises from: (i) SMOCU's global, non-myopic utility formulation, which accounts for the expected uncertainty reduction impact of each query; and (ii) e-straddle's adaptive balancing of exploration and exploitation, modulated by the transformation function $g(\cdot)$.

Quantitative Comparison. Table 2 reports the final Bayesian error after 80 iterations for different transformation functions $g(\cdot)$ and exploration weights β in the *e-straddle* utility. For the continuous-query formulation, all transformations outperform the LP and LS baselines, with the best performance achieved by $g(r) = \sqrt{r}$ when $\beta = 2$. This result highlights the benefit of flexible reward shaping in modulating exploration–exploitation behavior. Table 3 summarizes the final Bayesian error for different initial sample sizes. In all configurations, the BALPI variants, particularly SMOCU and e-straddle (c), consistently achieve the lowest errors, demonstrating the robustness of the framework with respect to initialization. Importantly, Table 3 also provides a direct quantitative comparison against the representative Bayesian baselines BALD and MES. Across all three initialization settings, SMOCU and e-straddle(c) achieve lower final Bayesian error than both BALD and MES, showing that the gains of BALPI are not limited to comparisons with LP/LS alone.

Summary of Findings. Across both material systems used as demonstration problems, the BALPI framework consistently achieves faster convergence and lower Bayesian error than existing active learning and non-Bayesian phase identification methods. SMOCU provides globally informed sampling and robust recovery of sparse phase regions,

View Article Online
DOI: 10.1039/D5DD00459D

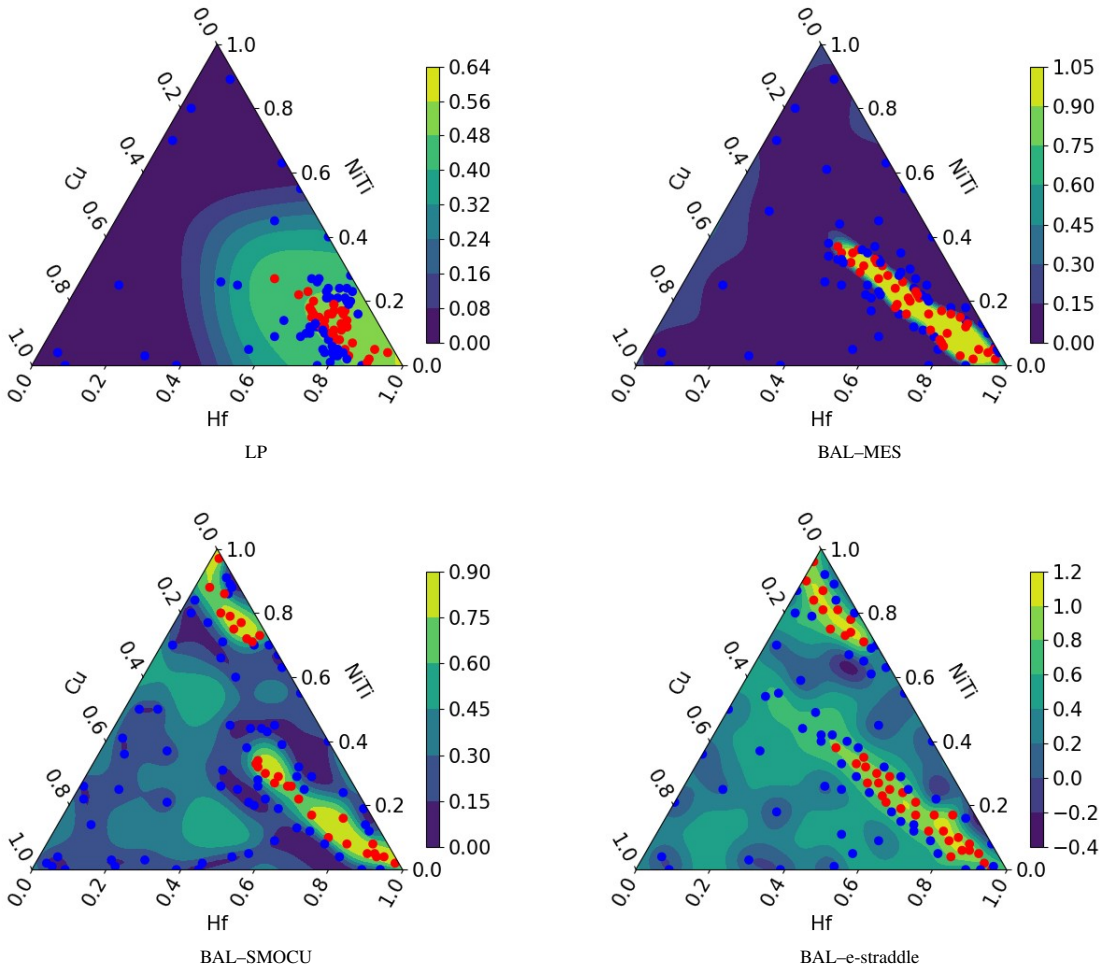

Figure 5: Final queried compositions and predicted BCC–B2 phase distributions after 80 iterations. Benchmark methods such as LP and MES tend to oversample low-uncertainty regions, missing smaller, disconnected BCC–B2 stability domains. In contrast, SMOCU- and e-straddle–guided BALPI efficiently identify multiple sparse, non-contiguous BCC–B2 regions, highlighting their ability to navigate complex phase landscapes with higher sampling efficiency.

while e-straddle offers an interpretable and computationally efficient alternative that performs especially well near complex phase boundaries. Together, these results demonstrate that BALPI constitutes a unified, uncertainty-aware, and computationally efficient approach for adaptive phase mapping in both discrete and continuous formulations.

4.4. Discussion

Interpretation of the Bayesian error metric. In this work, we adopt the Bayesian error as a unified performance metric for both classification and level-set formulations. Unlike standard accuracy or mean squared error, which quantify deviation from ground truth labels in a deterministic sense, the Bayesian error reflects *expected probability of misclassification* under the predictive posterior of the model. Formally, it is defined as the integral over the input space of the uncertainty-weighted misprediction of the model, as shown in Equation (11). This probabilistic interpretation is particularly well suited for Bayesian active learning, where the objective is to minimize predictive uncertainty as efficiently as possible—rather than to simply maximize accuracy on a fixed test set. A lower Bayesian error thus indicates that the model achieves not only greater predictive accuracy but also higher calibrated confidence with

View Article Online
DOI: 10.1039/D5DD00459D

	LP	LS	BALD	MES	SMOCU	e-straddle(b)	e-straddle(c)
4 samples	0.107±0.008	0.108±0.010	0.107±0.006	0.060±0.009	0.058±0.022	**0.050**±0.010	**0.045**±0.014
9 samples	0.097±0.010	0.096±0.009	0.103±0.009	0.060±0.013	**0.048**±0.007	0.056±0.010	**0.041**±0.004
16 samples	0.098±0.013	0.098±0.005	0.103±0.013	0.045±0.007	**0.041**±0.003	0.046±0.015	**0.038**±0.010

Table 3: Final Bayesian error after 80 iterations with varying numbers of randomly selected initial samples. SMOCU and *e-straddle(c)* achieve the lowest errors across all configurations.

respect to the ground-truth phase distribution. Importantly, the Bayesian error can be computed directly from the model's predictive distributions without requiring a separate validation set, making it an interpretable, uncertainty-aware, and model-agnostic evaluation criterion.

Practical benefits of Bayesian active learning for CALPHAD-guided discovery. A central motivation of this work is to demonstrate that uncertainty-aware sampling can dramatically reduce the computational effort required to map phase stability regions from CALPHAD simulations. Exhaustive CALPHAD computations over dense ternary or quaternary grids are typically infeasible due to the combinatorial growth of sampling points and the time cost of each equilibrium calculation. For instance, generating a full pseudo-ternary grid with 100×100 points may require several hours to days of computation, even when parallelized. By contrast, our BALPI framework converges to a comparable or superior phase boundary description using only tens to a few hundred adaptively selected samples—when considering an experimental campaign, the sample efficiency of BALPI relative to competing strategies is much more significant. BALPI's sample efficiency arises from the ability of the Bayesian utility functions—such as SMOCU and e-straddle—to automatically focus new queries near high-uncertainty regions, effectively balancing exploration of poorly sampled areas with exploitation near phase boundaries. The resulting acceleration makes BALPI well-suited not only for CALPHAD-based studies but also for integration into experimental design loops where querying the phase stability domain of target materials involves significant time or material cost.

Strengths and limitations. The results on the glass-ceramic (SiO_2–Al_2O_3–MgO) and shape-memory alloy (Ni–Ti–Hf–Cu) systems demonstrate that BALPI can capture disconnected phase stability regions and avoid overconfidence in sparsely sampled domains. The modularity of the framework enables flexible combinations of predictive models (GPC or GPR) and acquisition functions (MES, BALD, SMOCU, e-straddle), making it extensible to a wide range of materials systems. However, several limitations should be noted. First, Gaussian processes are most computationally efficient in low-dimensional spaces; direct scaling to systems beyond quaternary composition spaces may require sparse GP or deep kernel learning variants. Second, the current implementation assumes that CALPHAD serves as a reliable oracle and does not explicitly handle uncertainties in the thermodynamic database itself. Extending BALPI to jointly account for CALPHAD parameter uncertainty or experimental measurement noise would provide a more comprehensive uncertainty model. Finally, while SMOCU and e-straddle balance exploration and exploitation effectively, their optimization relies on empirical hyperparameters (k, β), and adaptive tuning of these parameters could further improve robustness across different systems.

Outlook. Despite these limitations, the present study establishes BALPI as a general and interpretable Bayesian active learning framework for phase diagram discovery. By unifying classification and level-set formulations under a probabilistic surrogate model, BALPI provides a natural interface between physics-based CALPHAD models and data-driven exploration strategies.

Looking ahead, the methodology presented here can be seamlessly extended to experimental workflows, enabling closed-loop autonomous materials design in which each new experiment is selected based on quantified uncertainty. In such settings, experimentally determined phase constitution states could be used not only for exploration, but also to refine existing CALPHAD models [35], including those constructed from computational approaches such as machine-learned interatomic potentials [36, 37]. Future work will focus on coupling BALPI with high-throughput synthesis and characterization platforms to enable such closed-loop integration.

View Article Online
DOI: 10.1039/D5DD00459D

While the present study focuses on binary phase identification to efficiently localize the stability region of a target phase, the BALPI framework is not intrinsically limited to two-phase settings. Extension to multi-phase phase-diagram construction is straightforward. For example, multi-class Gaussian process classification can be employed to directly model multiple phase labels, or, alternatively, parallel level-set estimation of individual phase-fraction fields can be used to track several phase boundaries simultaneously. In either formulation, BALPI retains its core objective: efficient localization of decision boundaries in thermodynamic state space.

Beyond identifying complete phase fields, one may instead target the zero phase fraction (ZPF) hyperplane [38] for phase i, which separates regions of state space where phase i is present from those where it is absent [2]. This formulation reduces the problem to a one-versus-all classification task, allowing BALPI to operate exactly as implemented here, without modifying the orignal algorithm in any way.

Importantly, the collection of ZPF hyperplane projections for all phases constitutes the two-dimensional projection of a phase diagram in a system of arbitrary dimensionality [38, 39]. Consequently, systematic identification of ZPF manifolds provides a scalable pathway toward high-dimensional phase-diagram reconstruction.

From a scalability perspective, extending BALPI to higher-dimensional composition or processing spaces will require surrogate models that can efficiently accommodate larger query budgets. While the present work employs standard Gaussian processes, future extensions will explore sparse and inducing-point GP formulations, batch acquisition strategies, and multi-fidelity active learning to balance computational tractability with uncertainty-aware exploration. These developments will enable BALPI to remain effective in quaternary and higher-order systems, where phase stability regions are often sparse, disconnected, and embedded in high-dimensional thermodynamic spaces.

5. Conclusion

In this work, we have introduced **BALPI**, a Bayesian active learning framework for adaptive phase diagram discovery grounded in CALPHAD thermodynamics. By unifying Gaussian Process Classification and Regression under a common probabilistic surrogate model, BALPI enables uncertainty-aware exploration through multiple acquisition strategies, including non-myopic formulations such as SMOCU and boundary-aware alternatives such as extended straddle. This dual capability—of supporting both discrete and continuous phase representations—makes BALPI a flexible and interpretable tool for phase mapping across a wide range of materials systems.

Demonstrated on two distinct thermodynamic systems—a ternary glass-ceramic and a quaternary shape memory alloy—BALPI consistently outperforms conventional sampling strategies, achieving faster convergence and lower Bayesian error while using a fraction of the CALPHAD query budget. Crucially, the framework succeeds even in scenarios with sparse and topologically fragmented stability regions, such as disconnected BCC-B2 pockets critical for the design and optimization of novel multi-component shape memory alloys. These results constitute strong evidence for the power of active, uncertainty-guided sampling grounded in Bayesian principles—not merely as a tool for efficient model training, but as an engine for discovery in systems where the ground truth occupies sparse, fragmented, or otherwise non-intuitive regions of chemical space.

Looking ahead, BALPI opens several promising research directions. First, its integration with high-throughput experimental workflows could enable closed-loop, autonomous discovery of phase boundaries, where surrogate uncertainty guides physical synthesis and characterization. Second, BALPI provides a natural interface for model refinement: experimentally observed phase constitutions could be used not just for exploration, but to iteratively update CALPHAD assessments—particularly those constructed from high-throughput DFT or MLIP data [37]. Third, embedding thermodynamic priors into the GP kernel—such as convexity constraints, symmetry relations, known limiting behaviors or topological rules of thermodynamic spaces—could improve sample efficiency and enhance generalization to unobserved regions of composition-temperature space.

Beyond phase diagrams, the BALPI paradigm is extensible to broader materials-informatics tasks involving phase

View Article Online
DOI: 10.1039/D5DD00459D

selection, stability prediction, and structural motif classification. In systems governed by sharp transitions or highly constrained free energy landscapes, the combination of interpretable surrogate models and targeted information acquisition offers a principled foundation for self-driving materials platforms. As phase mapping increasingly becomes a bottleneck in alloy design and microstructure engineering, frameworks like BALPI will be central to accelerating both understanding and deployment of next-generation materials.

Declaration of Competing Interest

The authors declare no competing interests.

Acknowledgements

The authors would like to acknowledge support from National Science Foundation (NSF) Grants No. DMREF-2119103, DMREF-1905325, CCF-1553281, SHF-2215573, and IIS-2212419. Research was also partly sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-22-2-0106 (HTMDEC-BIRDSHOT program). Computing resources at Texas A&M University were used for the calculations.

Authors Contributions

Fan: Conceptualization, Investigation, Formal Analysis, Methodology, Writing - Original Draft. **Wang**: Investigation, Formal Analysis, Writing - Review & Editing. **Vazquez**: Investigation, Writing - Original Draft. **Zhou**: Investigation, Formal Analysis, Methodology. **Karaman**: Conceptualization, Funding acquisition. **Arróyave**: Conceptualization, Supervision, Writing - Review & Editing, Funding acquisition. **Qian**: Conceptualization, Supervision, Writing - Review & Editing, Funding acquisition.

Data Availability

The code developed in this paper is made open accessible at the following GitHub repository: https://github.com/MingzhouFan97/BALPI. The associated dataset is publicly archived on Zenodo at https://doi.org/10.5281/zenodo.19896908.

References

[1] R. DeHoff, Thermodynamics in materials science, CRC Press, 2006.
[2] R. Arróyave, Phase stability through machine learning, Journal of Phase Equilibria and Diffusion 43 (6) (2022) 606–628.
[3] L. Kaufman, H. Bernstein, Computer calculation of phase diagrams. with special reference to refractory metals (1970).
[4] U. R. Kattner, The calphad method and its role in material and process development, Tecnologia em metalurgia, materiais e mineracao 13 (1) (2016) 3.
[5] P. E. Turchi, I. A. Abrikosov, B. Burton, S. G. Fries, G. Grimvall, L. Kaufman, P. Korzhavyi, V. R. Manga, M. Ohno, A. Pisch, et al., Interface between quantum-mechanical-based approaches, experiments, and calphad methodology, Calphad 31 (1) (2007) 4–27.
[6] P. Ren, Y. Xiao, X. Chang, P.-Y. Huang, Z. Li, B. B. Gupta, X. Chen, X. Wang, A survey of deep active learning, ACM computing surveys (CSUR) 54 (9) (2021) 1–40.
[7] S. Ament, M. Amsler, D. R. Sutherland, M.-C. Chang, D. Guevarra, A. B. Connolly, J. M. Gregoire, M. O. Thompson, C. P. Gomes, R. B. Van Dover, Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams, Science Advances 7 (51) (2021) eabg4930.
[8] C. Dai, S. C. Glotzer, Efficient phase diagram sampling by active learning, The Journal of Physical Chemistry B 124 (7) (2020) 1275–1284.

View Article Online
DOI: 10.1039/D5DD00459D

[9] D. O. Abranches, E. J. Maginn, Y. J. Colón, Activity coefficient acquisition with thermodynamics-informed active learning for phase diagram construction, AIChE Journal 69 (8) (2023) e18141.
[10] G. Deffrennes, K. Terayama, T. Abe, R. Tamura, A machine learning–based classification approach for phase diagram prediction, Materials & Design 215 (2022) 110497.
[11] K. Terayama, R. Tamura, Y. Nose, H. Hiramatsu, H. Hosono, Y. Okuno, K. Tsuda, Efficient construction method for phase diagrams using uncertainty sampling, Physical Review Materials 3 (3) (2019) 033802.
[12] Y. Xue, J. Bai, R. Le Bras, R. Bernstein, J. Bjorck, L. L. L. Longpre, S. Suram, R. Van Dover, J. Gregoire, C. Gomes, Phase-mapper: an ai platform to accelerate high throughput materials discovery, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 31, 2017, pp. 4635–4642.
[13] J. Bai, Y. Xue, J. Bjorck, R. Le Bras, B. Rappazzo, R. Bernstein, S. K. Suram, R. B. Van Dover, J. M. Gregoire, C. P. Gomes, Phase mapper: Accelerating materials discovery with ai, Ai Magazine 39 (1) (2018) 15–26.
[14] C. E. Rasmussen, H. Nickisch, Gaussian processes for machine learning (gpml) toolbox, The Journal of Machine Learning Research 11 (2010) 3011–3015.
[15] J. Deely, D. Lindley, Bayes empirical bayes, Journal of the American Statistical Association 76 (376) (1981) 833–841.
[16] C. Hardcastle, R. O'Mullan, R. Arroyave, B. Vela, Physics-informed gaussian process classification for constraint-aware alloy design, Digital Discovery (2025).
[17] T. Lookman, P. V. Balachandran, D. Xue, R. Yuan, Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design, npj Computational Materials 5 (1) (2019) 21.
[18] G. Zhao, E. Dougherty, B.-J. Yoon, F. J. Alexander, X. Qian, Bayesian active learning by soft mean objective cost of uncertainty, in: International Conference on Artificial Intelligence and Statistics, PMLR, 2021, pp. 3970–3978.
[19] G. Zhao, E. Dougherty, B.-J. Yoon, F. Alexander, X. Qian, Efficient active learning for gaussian process classification by error reduction, Advances in Neural Information Processing Systems 34 (2021) 9734–9746.
[20] R. Jiang, H. Zhou, H. Wang, S. S. Ge, Maximum entropy searching, CAAI Transactions on Intelligence Technology 4 (1) (2019) 1–8.
[21] Y. Gal, R. Islam, Z. Ghahramani, Deep bayesian active learning with image data, in: International conference on machine learning, PMLR, 2017, pp. 1183–1192.
[22] A. Gotovos, Active learning for level set estimation, Master's thesis, Eidgenössische Technische Hochschule Zürich, Department of Computer Science, (2013).
[23] P. I. Frazier, A tutorial on bayesian optimization, arXiv preprint arXiv:1807.02811 (2018).
[24] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner, W. Penny, Variational free energy and the laplace approximation, Neuroimage 34 (1) (2007) 220–234.
[25] H.-C. Kim, Z. Ghahramani, Bayesian gaussian process classification with the em-ep algorithm, IEEE Transactions on Pattern Analysis and Machine Intelligence 28 (12) (2006) 1948–1959.
[26] B. Bryan, R. C. Nichol, C. R. Genovese, J. Schneider, C. J. Miller, L. Wasserman, Active learning for identifying function threshold boundaries, Advances in neural information processing systems 18 (2005).
[27] S. A. Ansar, S. Bhattacharya, S. Dutta, S. Ghosh, S. Mukhopadhyay, Development of mullite and spinel coatings on graphite for improved water-wettability and oxidation resistance, Ceramics International 36 (6) (2010) 1837–1844.
[28] M. Leśniak, J. Partyka, K. Pasiut, M. Sitarz, Microstructure study of opaque glazes from sio2–al2o3–mgo–k2o–na2o system by variable molar ratio of sio2/al2o3 by ftir and raman spectroscopy, Journal of Molecular Structure 1126 (2016) 240–250.
[29] H. Karaca, E. Acar, H. Tobe, S. Saghaian, Nitihf-based shape memory alloys, Materials Science and Technology 30 (13) (2014) 1530–1544.
[30] H. E. Karaca, E. Acar, G. Ded, S. Saghaian, B. Basaran, H. Tobe, M. Kok, H. Maier, R. Noebe, Y. Chumlyakov, Microstructure and transformation related behaviors of a ni45. 3ti29. 7hf20cu5 high temperature shape memory alloy, Materials Science and Engineering: A 627 (2015) 82–94.
[31] J. Broucek, D. Khatamsaz, C. Cakirhan, S. H. Zadeh, M. Fan, G. Vazquez, K. Atli, X. Qian, R. Arroyave, I. Karaman, Design of high-temperature niticuhf shape memory alloys with minimum thermal hysteresis using bayesian optimization, Acta Materialia 286 (2025) 120651.
[32] J. Liu, L. Zhu, X. Huang, G. Cai, Z. Jin, Investigation of the phase equilibria in ti-ni-hf system using diffusion triples and equilibrated alloys, Calphad 58 (2017) 160–168.
[33] L. Lv, Z. Deng, H. Zhang, L. Liu, L. Zhang, Thermodynamic assessment of the ternary ni–ti-hf system, Calphad 81 (2023) 102546.
[34] K. Li, W. Zheng, Thermodynamic modeling of the ni–ti–hf system, Calphad 81 (2023) 102558.
[35] C. Kunselman, S. Zhu, D. Sariturk, R. Arroyave, Construction and tuning of calphad models using machine-learned interatomic potentials and experimental data: A case study of the pt-w system, arXiv preprint arXiv:2508.01028 (2025).
[36] S. Zhu, D. Sarıtürk, R. Arróyave, Accelerating calphad-based phase diagram predictions in complex alloys using universal machine learning potentials: Opportunities and challenges, Acta Materialia (2025) 120747.
[37] S. Zhu, D. Sariturk, R. Arroyave, Machine learning potentials for alloys: A detailed workflow to predict phase diagrams and benchmark accuracy, arXiv preprint arXiv:2506.16771 (2025).
[38] H. Gupta, J. Morral, H. Nowotny, Constructing multicomponent phase diagrams by overlapping zpf lines, Scripta metallurgica 20 (ISSN: 0036-9748) (1986).
[39] J. Morral, H. Gupta, A figure of merit for predicted phase diagrams, Journal of phase equilibria 13 (4) (1992) 373–376.

View Article Online
DOI: 10.1039/D5DD00459D

Data availability

The code developed in this paper is made open accessible at the following GitHub repository: https://github.com/MingzhouFan97/BALPI. The associated dataset is publicly archived on Zenodo at https://doi.org/10.5281/zenodo.19896908.