






Cite this: DOI: 10.1039/d6dd00009f

Jeweler-in-the-loop: personalized alloy color optimization *via* preference-based BO

Chase Katz, * Ting-Yu Yang,  Parker King, Md Shafiqul Islam, Brent Vela 
and Raymundo Arróyave

Many materials design attributes that are central to adoption, such as aesthetics, perceived quality, or user-specific preferences, are difficult to quantify directly, making preference feedback a practical proxy for optimization. Here we present a preference-driven Bayesian optimization framework and demonstrate it using ring color as a concrete case study. This case study simulates a scenario in which a jeweler synthesizes a ring with a given chemistry and presents it to a stakeholder, who expresses a preference relative to an incumbent ring color. Our approach, preferential Bayesian optimization (PBO), learns a latent utility function over alloy composition space, with perceived color obtained *via* a Thermo-Calc optical forward model. Using preference feedback, the framework iteratively proposes new alloy chemistries predicted to better align with user aesthetic preferences. After a limited number of proposed alloys and their associated colors, the user selects the most desirable option, guiding the search toward optimal aesthetic outcomes. We then evaluate the cost of the proposed alloys and identify a cost-aesthetic Pareto front, enabling informed trade-offs between affordability and visual appeal. The proposed framework is readily applicable to materials design problems in which subjective or hard-to-measure attributes play a dominant role in decision-making.

Received 9th January 2026
Accepted 2nd March 2026

DOI: 10.1039/d6dd00009f

rsc.li/digitaldiscovery

1. Introduction

Rings have long served as powerful symbols across cultures, embodying ideas such as commitment, heritage, status, and personal identity.^{1–3} The art of metallurgy, dating back over 6000 years, enabled ancient civilizations to work copper, gold, and bronze into functional and symbolic objects, often signifying wealth, power, or spiritual meaning.^{3,4} Today, the jewelry industry continues to thrive, with the global market valued at approximately \$366.8 billion in 2024 and projected to grow to \$578.5 billion by 2033, driven by rising disposable incomes, evolving fashion trends, and increasing demand for personalized accessories.⁵

This scale presents a significant financial opportunity for jewelers. In an ideal scenario with unlimited resources, a custom jeweler could fabricate every possible alloy within a given chemical design space and present the full palette of ring colors to the user for selection. However, exhaustive fabrication to explore alloy compositions for a preferred color is costly and time-consuming. The expense is particularly acute because precious metals such as gold and platinum can cost over \$130 and \$50 per gram,^{6,7} respectively, meaning that even a handful of prototype rings represents a substantial investment.

Traditional alloy design for rings and other jewelry has historically relied on trial-and-error synthesis and subjective evaluation of properties such as color. Moreover, aesthetic preferences are inherently personal and difficult to quantify, making systematic optimization of ring color for individual users challenging.⁸ Consequently, there is a need for methods that can efficiently identify optimal aesthetic outcomes while minimizing the number of prototypes required.

More broadly, many materials design problems are ill-posed as the optimization of a fully specified scalar or even a fixed multi-objective function. Key attributes that affect ring preference (aesthetics, tactile feel, surface finish, perceived quality, or designer intent) are not directly measurable, vary across stakeholders, and evolve as candidates are explored. In these settings, the objective is better treated as a latent preference model rather than a known function. Stakeholders can often provide reliable pairwise judgments (A vs. B) even when absolute scoring is inconsistent, making preference data a natural feedback channel for learning and optimization. In this context, “jeweler-in-the-loop” refers to a human-in-the-loop design workflow in which a jeweler or designer provides iterative pairwise preference judgments over candidate ring colors, guiding the optimization process without explicitly specifying a quantitative objective function.

Preference-based optimization formalizes this process by combining a latent-utility model with an acquisition strategy that selects the next comparison to maximize expected utility

Department of Materials Science and Engineering, Texas A&M University, College Station, TX 77843, USA. E-mail: ichasekatz@icloud.com



improvement or information gain.⁹ This yields an iterative loop that can efficiently explore high-dimensional composition spaces while remaining faithful to subjective or context-dependent criteria. The approach is compatible with conventional objectives and constraints, enabling joint optimization (e.g., preference-informed utility vs. cost, processability, or compliance) within a multi-objective PBO framework. In this work, alloy color is a concrete example of a hard-to-quantify attribute, but the method is intended as a general recipe for materials design problems where preference signals are the most reliable data available.

In this work, we present a preference-driven Bayesian optimization framework to identify ring colors that best match an individual user's aesthetic preferences.^{8,9} Preferential Bayesian optimization learns a latent utility function over alloy composition space, with color induced by the Thermo-Calc optical forward model, and proposes new alloy chemistries predicted to align with the user. We implement this as an iterative loop in which each query renders two candidate alloy colors and elicits a pairwise preference, using that feedback to update the utility model and drive subsequent proposals.

To demonstrate the proposed preference-based Bayesian optimization method, we perform an *in silico* study using Thermo-Calc's optical properties module within the Noble Metal Alloys Database (TCNOBL4) as the ground truth.^{10,11} Each query is analogous to fabricating a physical prototype. Each query retrieves a predicted color from Thermo-Calc's optical Model. The user is then presented with two candidate designs and selects the preferred option. After a fixed number of iterations, the resulting set of candidate alloys is analyzed in a multi-objective setting to identify the aesthetic-cost trade-off frontier, which serves as a decision aid for selecting a composition that balances visual appeal with affordability.

Preference-based Bayesian optimization has been widely studied, including multi-objective and scalable variants.^{8,9,12–17} Prior work has addressed preferences over explicit objectives, such as preference-order constraints within multi-objective BO frameworks that bias sampling toward preferred regions of the Pareto front, as well as extensions of preferential BO to multi-objective settings by scalarized Thompson sampling. Other studies have focused on algorithmic improvements, including preference exploration strategies for efficient multi-outcome optimization, decision-theoretic acquisition functions such as qEUBO, and improved inference schemes for preferential Gaussian processes that account for posterior skewness. Applications of preferential BO to subjective design problems, such as additive manufacturing parameter tuning, have also been demonstrated.

These prior works primarily focus on developing new acquisition strategies, inference schemes, or scalable variants of preferential Bayesian optimization. In contrast, this work does not introduce a new preferential BO algorithm. Instead, we adopt the established pairwise Gaussian-process preference model and deploy it within a materials-design setting, demonstrating how preference-based BO can be embedded within alloy composition design using a CALPHAD-based optical model.

Prior work has addressed preferences over explicit objectives, such as preference-order constraints within multi-objective BO frameworks that bias sampling toward preferred regions of the Pareto front,¹³ as well as extensions of preferential BO to multi-objective settings by scalarized Thompson sampling.¹² Other studies have focused on algorithmic improvements, including preference exploration strategies for efficient multi-outcome optimization,¹⁵ decision-theoretic acquisition functions such as qEUBO,¹⁷ and improved inference schemes for preferential Gaussian processes that account for posterior skewness.¹⁶ Applications to subjective design problems, such as additive manufacturing parameter tuning, have also been demonstrated.¹⁴

In contrast, this work does not introduce a new preferential BO algorithm. Instead, we adopt the established pairwise Gaussian-process preference model and deploy it within a materials-design setting, demonstrating how preference-based BO can be embedded within alloy composition design using a CALPHAD-based optical model.

The novelty of this work lies in the application and integration of established preference-based optimization within a materials-design workflow. To our knowledge, this is one of the first applied design studies to deploy Thermo-Calc's optical properties module within TCNOBL4 for alloy color prediction in an iterative design setting. This deployment enables a preference-driven, human-in-the-loop Bayesian optimization workflow for ring design that integrates a physically grounded optical forward model, human preference feedback, and cost trade-off analysis within alloy composition space. Ring color serves as a concrete demonstration of how preference-based optimization can be used in materials problems where key attributes are subjective or difficult to quantify directly. Together, these contributions establish a practical use case for the Thermo-Calc optical model and a generalizable template for preference-based materials design.

2. Methods

2.1. Modeling latent color-preferences: pairwise Gaussian processes

We summarize the standard pairwise Gaussian-process preference learning formulation for completeness and to make the manuscript self-contained; this work adopts the established model without modification and focuses on its application to materials design.^{8,9}

We consider a preference-learning setting in which each candidate alloy is represented by a composition vector $x \in \mathbb{R}^d$, where each component of x corresponds to the fractional concentration of a constituent element e in the alloy design space. There are $d + 1$ elements in the alloy system, with x constrained to the composition simplex (non-negative fractions summing to a fixed total). The goal is to infer a latent utility function $f(x)$ from noisy pairwise ring comparisons.⁹ The training set consists of a collection of items $X = \{x_1, \dots, x_n\}$ together with m observed preferences, where the k -th preference states that x_{i_k} is preferred to x_{j_k} .



Pairwise comparisons are encoded through a design matrix $D \in \mathbb{R}^{m \times n}$, where each row corresponds to a single incumbent–challenger preference (i_k, j_k) and contains a +1 in the column of the preferred item and a –1 in the column of the non-preferred item, with zeros elsewhere, as shown in Table 1. Let $f = (f_1, \dots, f_n)^\top$ denote the latent utilities at the n candidate alloys, with $f_i = f(x_i)$. The design matrix maps utilities to comparison differences *via*

$$(Df)_k = f_{i_k} - f_{j_k}, \quad \text{so that } Df = \begin{bmatrix} f_{i_1} - f_{j_1} \\ f_{i_2} - f_{j_2} \\ \vdots \\ f_{i_m} - f_{j_m} \end{bmatrix}.$$

We use a Gaussian process (GP) prior to model the latent utility:

$$f \sim \mathcal{GP}(m(x), k_\theta(x, x')).$$

The prior mean $m(x)$ is taken to be constant, providing an uninformative, non-biasing baseline. The covariance function is chosen as a scaled radial basis function kernel,

$$k_\theta(x, x') = s^2 \exp\left(-\frac{1}{2} \sum_{e=1}^d \frac{(x_e - x'_e)^2}{\ell_e^2}\right),$$

where $\ell_e > 0$ are lengthscale parameters (optionally distinct across dimensions) and $s^2 > 0$ is an output scale. Kernel hyperparameters were restricted to positive values through the default parameterization and constraints in GPyTorch/BoTorch. The multi-objective model used the BoTorch PairwiseGP, which applies default weakly informative priors to the kernel lengthscale and output scale (the prior distributions and parameters are specified in the SI). The single-objective study used a GPyTorch RBF kernel with a ScaleKernel and no explicit hyperparameter priors. Alloy compositions were represented as normalized fractions in $[0, 1]$, and no manual hyperparameter tuning was performed. Implementation details and hyperparameter settings are described in the SI and documented in the repository.

Observed comparisons are modeled through a probit likelihood.⁹ Conditional on latent utilities at the compared items, the probability that x_{i_k} is preferred to x_{j_k} is

$$p(x_{i_k} > x_{j_k} | f) = \Phi(f_{i_k} - f_{j_k}),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Assuming independent comparisons given f , the log likelihood is

$$\log p(C|f) = \sum_{k=1}^m \log \Phi((Df)_k).$$

Because the probit likelihood yields a non-Gaussian posterior, inference proceeds *via* a Laplace approximation centered at the maximum *a posteriori* (MAP) utilities f_{MAP} .⁹ Define the negative log posterior,

$$S(f) = -\log p(C|f) + \frac{1}{2}(f - m)^\top K^{-1}(f - m),$$

where m is the prior mean vector evaluated at X and K is the kernel matrix $K_{ij} = k_\theta(x_i, x_j)$. The MAP estimate satisfies $\nabla S(f) = 0$ and is obtained using Newton–Raphson iterations that exploit analytic expressions for the gradient and Hessian of $S(f)$. The Hessian takes the form

$$\nabla^2 S(f) = K^{-1} + W,$$

where W is the negative Hessian of the log likelihood evaluated at the current utilities. The resulting Laplace posterior is a Gaussian approximation,

$$p(f|C, \theta) \approx \mathcal{N}(f_{\text{MAP}}, (K^{-1} + W)^{-1}).$$

To learn the kernel hyperparameters $\theta = \{s^2, \ell_1, \dots, \ell_d\}$, we maximize the Laplace-approximated marginal log likelihood (evidence).⁹ Because the probit likelihood renders the exact evidence intractable, we expand the log posterior around the MAP estimate f_{MAP} . Let $W = -\nabla^2 \log p(C|f)|_{f_{\text{MAP}}}$ denote the negative Hessian of the log likelihood evaluated at the MAP point, and let K be the prior covariance matrix. The Laplace approximation yields

$$\begin{aligned} \log p(C|\theta) &\approx \log p(C|f_{\text{MAP}}) - \frac{1}{2}(f_{\text{MAP}} - m)^\top K^{-1}(f_{\text{MAP}} - m) \\ &\quad - \frac{1}{2} \log |K^{-1} + W| + \text{const.} \end{aligned}$$

This expression balances data fit at the MAP utilities, the GP prior regularization, and the Occam factor $\log |K^{-1} + W|$, which penalizes models that are unnecessarily flexible. Maximizing this approximate marginal likelihood provides a principled and tractable way to learn the hyperparameters for pairwise-probit Gaussian process preference models.

Table 1 Illustrative design matrix D for pairwise comparisons in the preference-learning model

	x_1	x_2	x_3	x_4	x_5	x_6
$(i_1, j_1) = (1, 2)$ incumbent wins	+1	-1	0	0	0	0
$(i_2, j_2) = (1, 3)$ challenger wins	-1	0	+1	0	0	0
$(i_3, j_3) = (3, 4)$ incumbent wins	0	0	+1	-1	0	0
$(i_4, j_4) = (3, 5)$ challenger wins	0	0	-1	0	+1	0
$(i_5, j_5) = (5, 6)$ incumbent wins	0	0	0	0	+1	-1



2.2. Prediction with pairwise Gaussian processes

For a set of test points X_* , let $k_* = k(X_*, X)$ and $K_{**} = k(X_*, X_*)$. Under the Laplace approximation, the posterior predictive distribution of $f_* = f(X_*)$ is Gaussian with mean

$$\mu_* = m(X_*) + k_* K^{-1} (f_{\text{MAP}} - m(X)),$$

and covariance

$$\Sigma_* = K_{**} - k_* (K^{-1} + W)^{-1} k_*^\top.$$

This form reduces to standard GP prediction when $W = 0$ and incorporates preference information through the curvature of the likelihood at the MAP point.

Given the Gaussian predictive distribution for the latent utilities, the model can evaluate the probability that one candidate is preferred over another. For two test compositions $a, b \in X_*$, let μ_a, μ_b denote their predictive means and let

$$\sigma_{ab}^2 = \Sigma_{*,aa} + \Sigma_{*,bb} - 2\Sigma_{*,ab}$$

be the predictive variance of their utility difference. Under the probit likelihood, the preference probability is

$$\Pr(a > b) = \Phi\left(\frac{\mu_a - \mu_b}{\sigma_{ab}}\right),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Thus, the utility posterior induces a predictive distribution over pairwise preferences. This quantity is used implicitly in the acquisition function, since the posterior means μ_* enter the surrogate model and the posterior covariances Σ_* govern uncertainty in the latent utility.

2.3. Acquisition function: expected improvement (single objective)

Because the pairwise GP provides a posterior distribution over the latent utility $f(x)$, this latent function can be optimized directly using standard acquisition criteria within a PBO framework. In the single-objective setting, we select new candidates by maximizing Expected Improvement (EI).¹⁸ We implemented the optimization loop in BoTorch.¹⁹

Let $f(x)|\mathcal{D}_t \sim \mathcal{N}(\mu_t(x), \sigma_t^2(x))$ denote the Laplace-approximated posterior at iteration t . Let

$$f_t^+ = \max_{x \in X_t} \mu_t(x)$$

be the current best predicted utility among evaluated designs. The improvement achieved by evaluating a new candidate x is

$$I(x) = \max(0, f(x) - f_t^+),$$

which is a random variable under the posterior. The Expected Improvement is the posterior expectation of $I(x)$,

$$\alpha_{\text{EI}}(x) = \mathbb{E}[I(x)|\mathcal{D}_t].$$

Because the posterior over $f(x)$ is Gaussian, EI admits a closed form. Define the standardized improvement

$$z(x) = \frac{\mu_t(x) - f_t^+}{\sigma_t(x)},$$

and denote by $\Phi(\cdot)$ and $\phi(\cdot)$ the standard normal cumulative distribution function and probability density function. Then

$$\alpha_{\text{EI}}(x) = (\mu_t(x) - f_t^+) \Phi(z(x)) + \sigma_t(x) \phi(z(x)).$$

In this formulation, the latent utility GP plays the same role as an ordinary regression GP: its posterior mean $\mu_t(x)$ reflects how promising a candidate is according to inferred preferences, while the posterior variance $\sigma_t^2(x)$ quantifies uncertainty due to limited or ambiguous comparisons. Thus EI naturally balances exploitation of compositions predicted to yield high latent utility with exploration of regions where uncertainty in $f(x)$ remains large, even though the model is trained only on pairwise preference data.

2.4. Acquisition function: expected area improvement (bi-objective)

In this work, we consider the design space evaluated according to two conflicting objectives: (i) a latent utility function $u(x)$ describing color preference (to be maximized), inferred through pairwise comparison experiments, and (ii) a cost function $c(x)$ (to be minimized).²⁰ Following standard practice, we convert cost to a maximization objective by working with the transformed value

$$f_2(x) = -c(x),$$

so that the two-objective vector is

$$f(x) = (f_1(x), f_2(x)) = (u(x), -c(x)).$$

Each objective is modeled using an independent GP. The GP posterior at any unobserved point x provides predictive distributions

$$f_1(x) \sim \mathcal{N}(\mu_1(x), \sigma_1^2(x)), \quad f_2(x) \sim \mathcal{N}(\mu_2(x), \sigma_2^2(x)),$$

where independence between the two models is assumed for analytical tractability.

Let \mathcal{P} denote the current set of non-dominated (Pareto optimal) points in the bi-objective space. We choose a reference point $r = (r_1, r_2)$ that is dominated by all feasible objective values. The dominated hypervolume (in this case, area) is defined as

$$\text{HV}(\mathcal{P}) = \lambda\left(\bigcup_{p \in \mathcal{P}} [p_1, r_1] \times [p_2, r_2]\right),$$

where $\lambda(\cdot)$ denotes Lebesgue measure (area).^{20,21}

If a new point $f(x)$ is added, the hypervolume improvement is

$$\text{HVI}(x) = \max\{\text{HV}(\mathcal{P} \cup \{f(x)\}) - \text{HV}(\mathcal{P}), 0\}.$$

Since $f_1(x)$ and $f_2(x)$ are uncertain under the GP posteriors, we evaluate the Expected Hypervolume Improvement (EHVI):²¹



$$\text{EHVI}(x) = \mathbb{E}[\text{HVI}(x)].$$

For the two-objective case, the Pareto frontier can be written as an ordered sequence of points

$$p^{(1)}, p^{(2)}, \dots, p^{(K)},$$

sorted such that $p_1^{(k)}$ increases and $p_2^{(k)}$ decreases. These points divide the objective space into disjoint “cells” A_1, \dots, A_{K+1} .²¹ Let

$$a_k = p_1^{(k)}, \quad b_k = p_2^{(k-1)},$$

with the conventions

$$p^{(0)} = (r_1, p_2^{(1)}), \quad p^{(K+1)} = (p_1^{(K)}, r_2).$$

A newly sampled point $f(x)$ contributes hypervolume improvement only if it falls inside one of these cells. In such a case, the improvement is exactly the area of the rectangle:

$$\text{HVI}(x) = (f_1(x) - a_k)(b_k - f_2(x)), \quad \text{for } f(x) \in A_k.$$

In this application, the first objective models a latent preference derived from pairwise comparisons over alloy colors, while the second objective captures the manufacturing cost. EHVI evaluates each candidate design x by the expected increase in the area dominated by the Pareto frontier. This leads to principled exploration of alloy compositions that jointly improve color aesthetics and reduce cost. Because the hypervolume measures the trade-off surface between objectives, maximizing EHVI efficiently pushes the frontier outward, ensuring consistent progress toward discovering superior alloy designs.

In our preference-based bi-objective Bayesian optimization framework, the acquisition function (EHVI) is defined over individual candidate points, while the latent utility model requires pairwise comparisons for learning. Thus, at each iteration, we first identify the candidate that maximizes the analytical expected hypervolume improvement, selecting

$$x_{n+1} = \underset{x \in \mathcal{X}}{\text{argmax}} \text{EHVI}(x).$$

Because the latent utility is inferred from pairwise feedback, this EHVI-selected challenger is not evaluated alone, but is compared against the current incumbent x^* , defined as the design with the highest posterior mean utility under the PairwiseGP model,

$$x^* = \underset{x_i \in \mathcal{D}_n}{\text{argmax}} \hat{u}(x_i).$$

The next experiment therefore consists of presenting the pair (x^*, x_{n+1}) to the user, who indicates which of the two is preferred. This comparison updates the latent utility posterior, while the cost of the challenger is evaluated directly. This incumbent–challenger mechanism allows EHVI to guide exploration in the multi-objective space, while remaining consistent with the structure of pairwise preference learning.

2.5. Ground truth for alloy color: Thermo-Calc optical model

Alloy color was computed using the Thermo-Calc optical properties model, which predicts visible appearance of metals from equilibrium phase constitution and first-principles-informed dielectric data.^{10,11} For each composition x at the specified thermodynamic condition, the workflow proceeds as follows.

The electronic-structure origins of the dielectric response for individual elements (*e.g.*, Pt or Al) are not analyzed explicitly here; instead, we treat the Thermo-Calc optical properties model as a ground-truth oracle that deterministically maps alloy composition to visible color for the purposes of preference-based optimization.

The equilibrium phases $\{\alpha_i\}_{i=1}^N$ present at a given alloy composition x and their corresponding volume fractions $\{v_i\}_{i=1}^N$ were obtained from the Equilibrium Calculator. Here, N denotes the number of thermodynamically stable phases at composition x , α_i labels the i -th equilibrium phase, and v_i is its volume fraction, satisfying $\sum_{i=1}^N v_i = 1$. These phase identities and volume fractions constitute the phase-resolved inputs required for subsequent optical mixing calculations.

For each phase α_i , the complex dielectric function $\varepsilon_i(\lambda) = \varepsilon_{1,i}(\lambda) + i\varepsilon_{2,i}(\lambda)$ was taken from the Thermo-Calc noble-metal optical database as a function of wavelength λ .^{11,22} For multi-phase alloys, these were combined using the generalized Bruggeman effective-medium model to obtain an effective dielectric function $\varepsilon_{\text{eff}}(\lambda)$ satisfying

$$\sum_{i=1}^N v_i \frac{\varepsilon_i(\lambda) - \varepsilon_{\text{eff}}(\lambda)}{\varepsilon_i(\lambda) + (N-1)\varepsilon_{\text{eff}}(\lambda)} = 0, \quad (1)$$

where v_i and ε_i are the volume fraction and dielectric function of phase α_i , respectively.^{23,24}

The effective complex refractive index $\tilde{n}(\lambda) = n(\lambda) + ik(\lambda)$ was computed from $\varepsilon_{\text{eff}}(\lambda)$ *via*

$$n(\lambda) = \sqrt{\frac{\sqrt{\varepsilon_1^2(\lambda) + \varepsilon_2^2(\lambda)} + \varepsilon_1(\lambda)}{2}}, \quad (2)$$

$$k(\lambda) = \sqrt{\frac{\sqrt{\varepsilon_1^2(\lambda) + \varepsilon_2^2(\lambda)} - \varepsilon_1(\lambda)}{2}}, \quad (3)$$

where ε_1 and ε_2 are the real and imaginary parts of ε_{eff} .

Assuming light incident from air onto a semi-infinite alloy surface, the wavelength-dependent reflectance $R(\lambda)$ was calculated using the Fresnel equations for s- and p-polarized light at incidence angle θ :

$$R_s(\lambda) = \left| \frac{\tilde{n}_2 \cos \theta - \tilde{n}_1 \cos \theta_t}{\tilde{n}_2 \cos \theta + \tilde{n}_1 \cos \theta_t} \right|^2, \quad (4)$$

$$R_p(\lambda) = \left| \frac{\tilde{n}_2 \cos \theta - \tilde{n}_1 \cos \theta_t}{\tilde{n}_2 \cos \theta + \tilde{n}_1 \cos \theta_t} \right|^2, \quad (5)$$

with $\tilde{n}_1 = 1$ (air), $\tilde{n}_2 = \tilde{n}(\lambda)$ (alloy), and θ_t determined by Snell's law $\tilde{n}_1 \sin \theta = \tilde{n}_2 \sin \theta_t$. The unpolarized reflectance was taken as $R(\lambda) = \frac{1}{2}(R_s(\lambda) + R_p(\lambda))$.²⁴



Visible color was obtained by integrating the reflectance spectrum over the visible range (400–700 nm) in the CIE XYZ framework:²⁵

$$X = k \int_{400}^{700} R(\lambda)S(\lambda)\bar{x}(\lambda)d\lambda, \quad (6)$$

$$Y = k \int_{400}^{700} R(\lambda)S(\lambda)\bar{y}(\lambda)d\lambda, \quad (7)$$

$$Z = k \int_{400}^{700} R(\lambda)S(\lambda)\bar{z}(\lambda)d\lambda, \quad (8)$$

where $S(\lambda)$ is the spectral power distribution of the chosen standard illuminant (default D65), \bar{x} , \bar{y} , \bar{z} are the CIE color-matching functions for the selected standard observer (2° or 10°), and k is a normalization constant such that $Y = 100$ for a perfectly reflecting surface.²⁵ The resulting tristimulus vector (X, Y, Z) was then converted to the reported color space (CIE Lab or sRGB) using standard CIE transforms.

2.6. Definition of design space

It is of interest to assess how many alloy compositions within this high-dimensional design space can be evaluated using the Thermo-Calc optical properties model. The alloy design space, denoted \mathcal{X} , was defined as a quinary composition manifold spanned by {Au, Ag, Cu, Pt, Al}, with five composition variables constrained to sum to 100 at%, and sampled on a uniform grid in 5 at% increments. Only quinary alloys were considered; that is, each candidate composition contained at least 5 at% of every element, which substantially increases the combinatorial complexity of the color-composition mapping and provides a stringent test of the Thermo-Calc noble-metal optical database. Enumeration of this constrained 5 at% grid yielded 3876 distinct compositions. For each $x \in \mathcal{X}$, the Thermo-Calc optical properties-noble model was used to compute equilibrium phase fractions and wavelength-resolved optical properties.¹¹ Due to database coverage limitations, complete dielectric and optical solutions were successfully obtained for 933 of the 3876 quinary compositions, as shown in Fig. 1. These 933 alloys constitute

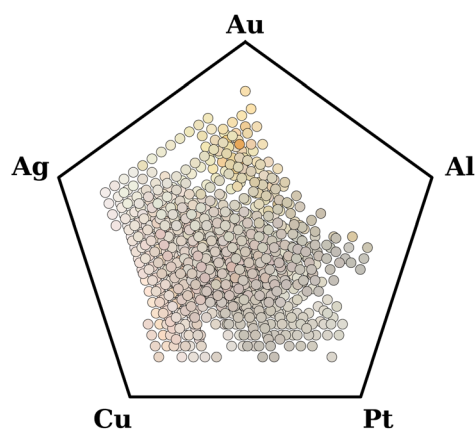


Fig. 1 Affine projection of calculated colors using Thermo-Calc's optical property module.

the design space that is to be explored with optimization in search of optimal color.

2.7. Preference-based Bayesian optimization of ring color

We iteratively learned an individual user's color preference over the quinary composition space using preference-based Bayesian optimization. Each alloy composition x induces a visible ring color *via* the optical forward model described above; in an ideal experimental workflow this color would be obtained from a synthesized ring, but here it is computed *in silico*. Preference learning is formulated through a latent scalar utility function $f(x)$ that represents the user's (unobserved) subjective valuation of ring color. The objective is to identify compositions that maximize $f(x)$ while minimizing the number of user queries.

At iteration t , the user is presented with two candidate rings having compositions $x_t^{(a)}$ and $x_t^{(b)}$. For each candidate, the corresponding color is predicted by the Thermo-Calc optical model and rendered onto a ring visualization to facilitate realistic comparison. The user provides a binary preference,

$$x_t^{(a)} > x_t^{(b)} \Leftrightarrow f(x_t^{(a)}) > f(x_t^{(b)}),$$

which is recorded as a noisy observation of the ordering of latent utilities. Collecting m_t such comparisons yields a dataset $\mathcal{D}_t = \{(x_i, x_j) : x_i > x_j\}_{k=1}^{m_t}$.

At each iteration, we select the next candidate by

$$x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}} \alpha_{\text{EI}}(x),$$

where \mathcal{X} is the candidate composition set. The newly selected composition is paired with a reference candidate (*e.g.*, the current incumbent best) to form a two-ring query shown to the user; the resulting preference is added to \mathcal{D}_t , the Pairwise GP posterior is updated, and the process repeats. In this way, user feedback on rendered ring colors is converted into pairwise constraints on $f(x)$, allowing the model to efficiently converge to compositions that best match the user's aesthetic preference.

3. Results and discussion

3.1. Optimizing color

In this case study, we conducted a color-preference-based ring optimization campaign for 25 iterations. In practice, a user would be presented with two color options and would pick the option they preferred. However, because color preference is inherently subjective and varies across individuals, we adopted an *in silico* benchmark to enable controlled evaluation. Specifically, to have a "simulated user" select its preferred ring colors, we defined a synthetic ground-truth latent preference function whose optimum corresponds to a canonical gold appearance, *i.e.*, the alloy (among the 933 candidates for which the optical model converged) with the highest Au content. This would be $\text{Ag}_{0.5}\text{Al}_{0.5}\text{Au}_{80}\text{Cu}_{0.5}\text{Pt}_{0.5}$. A "gold likeness" of a given color is assigned by using an equally weighted combination of color distance in HSV space and difference in a hue-based warmth metric between the candidate and the optimum color. This allows the model to have a quantitative value for a typically qualitative measurement as a substitute for a real person. The



simulated user was not assumed to know this preference *a priori*; instead, each query returned a pairwise choice consistent with this latent “gold” utility, thereby emulating an internal but undisclosed aesthetic target.

Starting from an initial set of candidate rings, the campaign proceeded by iteratively proposing two color-rendered options, eliciting a preference, and updating the pairwise Gaussian process posterior over latent utility. Over 25 iterations, the model progressively concentrated sampling near compositions predicted to yield higher gold-likeness, producing a clear trajectory toward the gold-dominant region of the design space, as seen in Fig. 2. To complement this visualization, the SI reports the optimization quantitatively, showing the R, G, and B channel values of sampled alloys moving toward those of the target gold color over successive iterations. This convergence demonstrates that preference-based Bayesian optimization can efficiently infer and exploit a user's latent aesthetic utility using only pairwise feedback.^{8,9}

These results indicate that the framework converges, as queried colors become increasingly gold-like over successive iterations. To assess the robustness of this behavior, we performed additional runs under varied hyperparameter and initialization settings; the resulting convergence trends were consistent and are reported in the SI. We further evaluated robustness to surrogate-model choices by repeating the optimization with alternative Gaussian-process kernels and a random sampling baseline, with results provided in the SI.

3.2. Optimizing color and cost

In the previous case study, the user's selections drove the preference-based optimization toward increasingly gold-like colors, reflecting the latent aesthetic utility learned during the campaign. However, in practical jewelry design, cost is a dominant and often competing consideration: Au is a particularly expensive alloying element, and if a user has a higher preference for gold colors, that usually entails higher material cost.

In this case study, we demonstrate a bi-objective optimization over ring color and price. As before, we model the color objective with a pairwise GP, while price is modeled with a standard GP regressor. Because price is relatively easy to

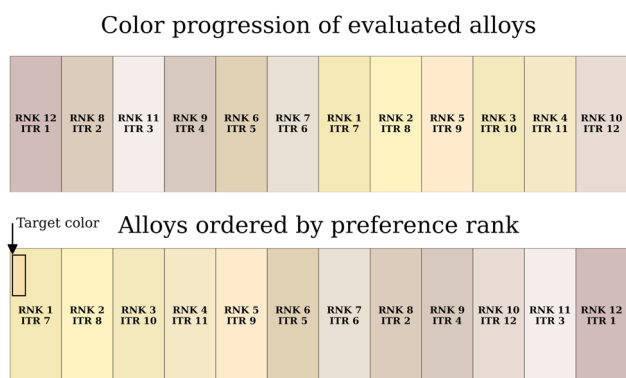
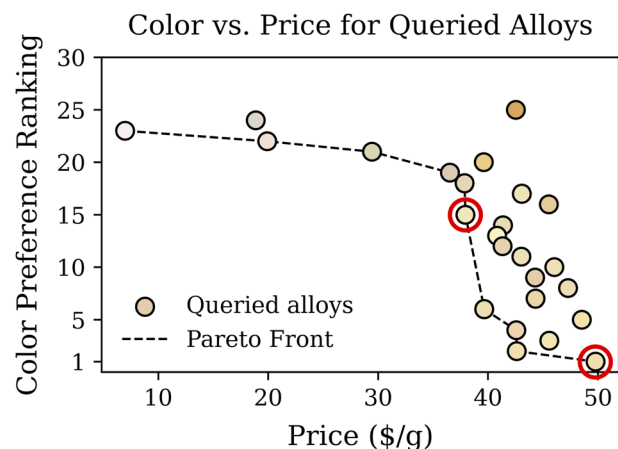


Fig. 2 Progress of evaluated alloy colors along with their associated preference ranking.



Preference Ranking	1	15
Ring Price	\$3683.15	\$2785.89
Money Saved	\$897.26	

Fig. 3 Pareto front of preference rank versus price per gram, showing the trade-off between a dummy user's learned preferred ring color and the cost of the ring.

predict from rule-of-mixtures estimates of elemental costs, we train the price GP once using all candidate alloys upfront; only the preference GP for color is updated iteratively as new comparisons are collected. Each iteration, the EHVI acquisition function then combines predictions from both GPs to propose new alloys that are expected to improve the current color-price Pareto front, *i.e.*, alloys whose colors the user is likely to prefer at an attractive price point. As before, we initialize the optimization scheme with 2 initial data points (a Pairwise GP needs to begin with 2 points as it needs an initial preference). We then run the optimization campaign for 25 iterations.

To analyze the results of this optimization campaign, we visualize the preference-rank-*versus*-price Pareto front, as shown in Fig. 3. As expected, more metallic silver-like colors tend to have lower costs, whereas more gold-like colors (*i.e.*, those closer to the preferred target color) are generally more expensive. A key advantage of this approach is that it presents users with a menu of Pareto-efficient options: they can select alloys along the color-price front according to their own willingness to pay for a given color when choosing candidates for further scale-up.²⁰

4. Conclusions

This work demonstrates a preference-based, human-in-the-loop Bayesian optimization framework for ring alloy design that unites materials science, statistical modeling, and jewelry customization. Using Thermo-Calc's optical model and noble-metal databases, we constructed a quinary Au–Ag–Cu–Pt–Al design space, computed equilibrium phase constitutions, and mapped each composition to a predicted visible color. These computations provide a physically grounded link between alloy chemistry and ring appearance.



On top of this forward model, we implemented PBO to infer a latent utility function over alloy composition from pairwise comparisons. In a simulated-user setting, the framework efficiently retrieved target colors within a limited number of queries, illustrating that pairwise feedback alone is sufficient to drive the search toward compositions that match an individual aesthetic preference. Extending this approach to a bi-objective setting, we incorporated a cost model and used expected hypervolume improvement to construct a color-cost Pareto front. This front provides users or designers with a transparent set of trade-off options between visual appeal and material expense.

Although the present study relies on *in silico* data and an idealized user model, the results highlight a practical path toward interactive, data-driven alloy design for jewelry. Future work will focus on experimental validation of selected compositions, integrating real user studies, and developing a more accessible graphical interface that connects preference learning, Thermo-Calc simulations, and visualization of color-cost trade-offs. Taken together, these developments point toward a customizable ring design workflow in which users can explore personalized aesthetic choices within physically realistic and economically informed alloy spaces.

Author contributions

Chase Katz: writing – original draft, writing – review & editing, visualization, software, investigation, formal analysis, data curation. Ting-Yu Yang: writing – original draft, writing – review & editing, visualization, software, investigation, formal analysis, data curation. Parker King: writing – original draft, writing – review & editing, visualization, software, investigation, formal analysis, data curation. Md Shafiqul Islam: writing – original draft, writing – review & editing, visualization, software, investigation, formal analysis, data curation, supervision. Brent Vela: writing – original draft, writing – review & editing, visualization, software, investigation, formal analysis, supervision. Raymundo Arróyave: writing – review & editing, supervision, funding acquisition.

Conflicts of interest

There are no conflicts of interest.

Data availability

The full implementation of the preferential Bayesian optimization color selection model is archived on Zenodo. The repository contains the complete source code and interactive notebook used in this study, including alloy color visualization, Gaussian process – based preference modeling, interactive user preference updates, and visualization of aesthetic – cost trade-offs for tested alloys.

The code supporting this article can be accessed at: <https://doi.org/10.5281/zenodo.18164808>.

Optical color outputs shown in this article were generated using Thermo-Calc with a licensed database that cannot be

redistributed publicly. To enable transparency and reuse, the workflow reports and visualizes obfuscated output data, included in the repository. Reproducing the full workflow requires access to Thermo-Calc.

Supplementary information (SI): additional figures, tables, computational details, model parameters, and supporting data related to the preference-based Bayesian optimization workflow and alloy color optimization results. See DOI: <https://doi.org/10.1039/d6dd00009f>.

Acknowledgements

We acknowledge the support from the U.S. Department of Energy (DOE) ARPA-E ULTIMATE Program through Project DE-AR0001427 and ARPA-E CHADWICK Program through Project DE-AR0001988. RA also acknowledges the Army Research Laboratory (ARL) for support through Cooperative Agreement Number W911NF-22-2-0106, as part of the High-throughput Materials Discovery for Extreme Conditions (HTMDEC) program as supported by the BIRDSHOT Center at Texas A&M University. BV acknowledges the support of NSF through Grant No. 1746932 (GRFP) and 1545403 (NRT-D3EM). The authors gratefully acknowledge Thermo-Calc Software AB for providing access to Thermo-Calc and associated optical-property databases used in this work.

References

- 1 M. Cifarelli, Adornment, identity, and authenticity: Ancient jewelry in and out of context, *Am. J. Archaeol.*, 2010, **114**(1), 1–9, DOI: [10.3764/ajaonline114.1.Cifarelli](https://doi.org/10.3764/ajaonline114.1.Cifarelli). URL <https://ajaonline.org/online-museum-review/368/>.
- 2 P. Kiernan and K.-P. Henz, Rings from the Forbidden Forest: The function and meaning of Roman trinket rings, *J. Roman Archaeol.*, 2023, **36**(1), 73–95, DOI: [10.1017/S1047759423000211](https://doi.org/10.1017/S1047759423000211).
- 3 M. Abram, A new look at the Mesopotamian rod and ring: Emblems of time and eternity, *Stud. Antiq.*, 2011, **10**(1), 15–36. <https://scholarsarchive.byu.edu/studiaantiqua/vol10/iss1/5/>.
- 4 M. Thoury, B. Mille, T. Séverin-Fabiani, L. Robbiola, M. Réfrégiers, J.-F. Jarrige and L. Bertrand, High spatial dynamics–photoluminescence imaging reveals the metallurgy of the earliest lost-wax cast object, *Nat. Commun.*, 2016, **7**, 13356, DOI: [10.1038/ncomms13356](https://doi.org/10.1038/ncomms13356).
- 5 Grand View Research, *Jewelry market size, share & trends analysis report*, 2024, <https://www.grandviewresearch.com/industry-analysis/jewelry-market>.
- 6 Bullion-Rates, Gold prices – US dollars (USD) – November 2025, accessed November 22, 2025, <https://www.bullion-rates.com/gold/USD/2025-11-history.htm>.
- 7 Bullion-Rates, Platinum prices – US dollars (USD) – November 2025, accessed November 22, 2025, <https://www.bullion-rates.com/platinum/USD/2025-11-history.htm>.
- 8 E. Brochu, V. M. Cora and N. de Freitas, A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical



- reinforcement learning, *arXiv*, 2010, preprint, arXiv:1012.2599, DOI: [10.48550/arXiv.1012.2599](https://doi.org/10.48550/arXiv.1012.2599).
- 9 W. Chu and Z. Ghahramani, Preference learning with Gaussian processes, *Proceedings of the 22nd international conference on Machine learning*, ACM, 2005.
- 10 J.-O. Andersson, *et al.*, Thermo-Calc and DICTRA, computational tools for materials science, *Calphad*, 2002, **26**, 273–312.
- 11 Thermo-Calc Software AB, Noble metal based alloys database, 2025, <https://thermocalc.com/products/databases/noble-metal-based-alloys/>.
- 12 R. Astudillo, K. Li, M. Tucker, C. X. Cheng, A. D. Ames and Y. Yue, Preferential multi-objective bayesian optimization, *arXiv*, 2024, preprint, arXiv:2406.14699, DOI: [10.48550/arXiv.2406.14699](https://doi.org/10.48550/arXiv.2406.14699).
- 13 M. Abdolshah, A. Shilton, S. Rana, S. Gupta and S. Venkatesh, Multi-objective bayesian optimisation with preferences over objectives, *arXiv*, 2019, preprint, arXiv:1902.04228, DOI: [10.48550/arXiv.1902.04228](https://doi.org/10.48550/arXiv.1902.04228).
- 14 J. R. Deneault, W. Kim, J. Kim, Y. Gu, J. Chang, B. Maruyama, J. I. Myung and M. A. Pitt, Preferential bayesian optimization improves the efficiency of printing objects with subjective qualities, *Digital Discovery*, 2025, **4**(3), 723–737, DOI: [10.1039/D4DD00320A](https://doi.org/10.1039/D4DD00320A).
- 15 Z. J. Lin, R. Astudillo, P. I. Frazier and E. Bakshy, Preference exploration for efficient bayesian optimization with multiple outcomes, *arXiv*, 2022, preprint, arXiv:2203.11382, DOI: [10.48550/arXiv.2203.11382](https://doi.org/10.48550/arXiv.2203.11382).
- 16 S. Takeno, M. Nomura and M. Karasuyama, Towards practical preferential bayesian optimization with skew gaussian processes, *arXiv*, 2023, preprint, arXiv:2302.01513, DOI: [10.48550/arXiv.2302.01513](https://doi.org/10.48550/arXiv.2302.01513).
- 17 R. Astudillo, Z. J. Lin, E. Bakshy and P. I. Frazier, geubo: A decision-theoretic acquisition function for preferential bayesian optimization, *arXiv*, 2023, preprint, arXiv:2303.15746, DOI: [10.48550/arXiv.2303.15746](https://doi.org/10.48550/arXiv.2303.15746).
- 18 D. R. Jones, M. Schonlau and W. J. Welch, Efficient global optimization of expensive black-box functions, *J. Global Optim.*, 1998, **13**, 455–492.
- 19 M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson and E. Bakshy, BoTorch: a framework for efficient Monte Carlo Bayesian optimization, *Adv. Neural Inf. Process. Syst.*, 2020, **33**, 21524–21538.
- 20 M. T. M. Emmerich and A. H. Deutz, A tutorial on multiobjective optimization: fundamentals and evolutionary methods, *Nat. Comput.*, 2018, **17**(3), 585–609, DOI: [10.1007/s11047-018-9685-y](https://doi.org/10.1007/s11047-018-9685-y).
- 21 M. Emmerich and J.-W. Klinkenberg, The computation of the expected hypervolume improvement in multiobjective optimization, *Operations Research Proceedings*, 2011.
- 22 P. B. Johnson and R. W. Christy, Optical constants of the noble metals, *Phys. Rev. B*, 1972, **6**, 4370–4379.
- 23 D. A. G. Bruggeman, Berechnung verschiedener physikalischer Konstanten von heterogenen Substanzen. I. Dielektrizitätskonstanten und Leitfähigkeiten der Mischkörper aus isotropen Substanzen, *Ann. Phys.*, 1935, **416**, 636–664.
- 24 C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*, Wiley, 1983.
- 25 International Commission on Illumination (CIE), *Colorimetry*, CIE Publication 15:2004, Commission Internationale de l'Eclairage, Vienna, Austria, 3rd edn, 2004.

