

Bayesian Optimisation of Photopolymer Formulations: Priors, Constraints, Mixed Variables and Experimental Efficiency

Using probabilistic optimisation without hiding state, feasibility or uncertainty

Juan Seguro

3Dresyns by Resyner Technologies S.L.

© 2026 Juan Seguro. All Rights Reserved.

Correspondence: info@3Dresyns.com · www.3Dresyns.com

ER-402 · Engineering Review · Version 0.1 · 7 October 2026

Abstract

Bayesian optimisation can reduce experimental burden when each formulation experiment is costly, but photopolymer development combines continuous concentrations, discrete chemistry choices, conditional variables, multiple objectives and hard manufacturability constraints. This review maps Bayesian optimisation to that mixed design space without treating the optimiser as a substitute for formulation science. It distinguishes prior knowledge from unsupported bias, model uncertainty from measurement uncertainty, feasibility from performance, and exploratory efficiency from qualification evidence. The central recommendation is to optimise within explicitly versioned state and constraint domains, preserve unsuccessful experiments, and stop when the remaining decision value is low rather than when a model merely appears numerically stable. Published materials-discovery and adaptive-manufacturing studies demonstrate sample-efficient closed-loop search, but direct transfer to photopolymer formulation requires state-resolved validation.

Keywords: Bayesian optimisation · photopolymer formulation · Gaussian process · mixed variables · constrained optimisation · multiobjective optimisation · experimental design

1. Engineering problem

Photopolymer formulation spaces are combinatorial. Continuous variables such as additive fraction, absorber level or solids loading interact with discrete choices such as oligomer family, initiator class or filler treatment. Process settings add further dimensions. Exhaustive factorial exploration quickly becomes infeasible, which makes Bayesian optimisation attractive because it chooses experiments sequentially using a probabilistic surrogate and an acquisition rule [1-3]. The engineering risk is to confuse sample efficiency with truth. A Bayesian optimiser can be highly efficient at learning the response produced by a biased method, a drifting instrument or an incorrectly bounded design space. The optimisation architecture therefore has to include measurement and feasibility models, not only a property surrogate.

2. Search-space representation

Continuous, ordinal, categorical and conditional variables should be represented according to their physical meaning. A categorical photoinitiator identity is not a continuous coordinate merely because it is encoded by an integer. A filler surface treatment may only be meaningful when the filler is present. A post-cure temperature may be conditional on a selected post-cure route. Variables should be scaled using defensible engineering ranges. The bounds are not algorithmic conveniences; they define where the model is authorised to propose experiments. Changes to raw-material identity, wavelength, machine, specimen geometry or test protocol can invalidate a previously trained response surface unless transfer is demonstrated.

3. Priors and prior data

Priors can encode known smoothness, monotonic tendencies or plausible length scales, but they should not encode beliefs that are unsupported by the current chemistry. Historical formulation data may be valuable only after checking state equivalence. Combining data from different printers, spectra, temperatures or post-cures without explicit covariates can create a deceptively dense dataset that violates the exchangeability assumed by the surrogate. An informative prior is useful when it expresses actual prior evidence and remains challengeable by new data. It becomes hazardous when it constrains the model so strongly that contradictory experiments are treated as noise.

Table 1. Variable types in a photopolymer Bayesian optimisation problem

Variable type	Example	Required treatment
Continuous	absorber fraction; solids loading	bounded scale with units
Categorical	initiator family; filler identity	categorical kernel/descriptor; no fake ordering

4. Objective definition and multiobjective decisions

Photopolymer development rarely has a single objective. Cure speed, viscosity, dimensional fidelity, modulus, toughness, optical transmission, conductivity, washability and ageing can conflict. Weighted scalarisation is convenient but embeds a policy decision into one number. Pareto-based or constrained approaches preserve more of the engineering trade space. Hard requirements should normally be constraints rather than negative weights. A formulation that fails a non-negotiable viscosity, stability or safety limit should not become optimal because another property is sufficiently high.

5. Feasibility and failed experiments

Failed or unmeasurable experiments are informative. Precipitation, phase separation, sedimentation, non-cure, excessive exotherm, recoating failure or instrument saturation define the boundary of the feasible domain. Deleting them biases the search toward an unrealistically smooth design space. A separate feasibility model can estimate whether a candidate can be prepared and measured. The

acquisition function can then combine expected performance with probability of feasibility. This is particularly important when formulation constraints are nonconvex.

6. Acquisition under uncertainty

Acquisition functions balance exploitation of promising regions against exploration of uncertain regions. The choice should follow the decision objective. Early in a campaign, exploration may reduce structural uncertainty; later, the value may shift toward discriminating among a small number of viable candidates. Batch acquisition additionally has to account for parallel experiments that have not yet returned results. Model uncertainty should not be interpreted as measurement uncertainty. Repeated experiments and reference controls are still needed to estimate experimental variation. A small posterior variance from a surrogate does not prove that the measurement system is accurate.

Variable type	Example	Required treatment
Conditional	dispersion time when filler is present	activate only in relevant branch
Process state	wavelength; dose; temperature	model explicitly or hold fixed
Constraint	viscosity; stability; cure/no-cure	separate feasibility or hard gate

Table 2. What uncertainty does and does not mean

Quantity	Meaning	Not equivalent to
Surrogate posterior uncertainty	uncertainty of model prediction under assumptions	measurement accuracy
Replicate dispersion	repeatability under stated method	model uncertainty
Domain distance	lack of nearby training evidence	proof of failure
Constraint probability	estimated feasibility	qualification acceptance

7. Mixed-variable and conditional structure

Formulation spaces often contain switches: a reactive diluent family determines which concentration bounds are meaningful; a filler type determines whether dispersion variables exist; a cure chemistry determines which spectral window is relevant. Hierarchical or mixed-variable optimisation is preferable to flattening these relationships into an artificial Euclidean space. When the search crosses chemistry families, the optimiser should be tested for whether it is interpolating within a family or extrapolating across a regime boundary. Regime transitions can be chemically abrupt even when numeric descriptors look smooth (see ER-311-ER-320 for dynamic-network regimes).

8. Stopping rules and experimental efficiency

A campaign should stop because the decision is sufficiently resolved, not because an arbitrary experiment count has been reached. Useful stopping conditions include low expected value of information, stable Pareto membership under uncertainty, repeated selection of equivalent candidates, or attainment of a predefined engineering target with confirmation. Adaptive materials studies demonstrate the ability to reduce search burden [2-4], but the reported efficiency is problem-specific. Photopolymer studies should report the baseline against which efficiency is claimed and should include failed and rejected trials in the accounting.

9. Independent confirmation and transfer

The selected candidate must be remade independently and tested with locked methods. Transfer to another printer, wavelength, batch or geometry requires a separate test because the response surface was learned under a specific state distribution. Bayesian optimisation does not remove the need for transfer qualification. When historical data are reused, the paper should identify which variables were held equivalent, which were harmonised, and which remain possible sources of domain shift.

10. Qualification decision boundary

Bayesian optimisation is most useful in photopolymer development when it is treated as a disciplined experimental-design layer: explicit variable types, explicit constraints, transparent priors, retained failures, calibrated uncertainty and independent confirmation. It is not a mechanism for converting sparse, heterogeneous measurements into universal formulation laws. The scientific contribution of an optimisation study should therefore be judged both by the candidate it finds and by the reproducibility and auditability of the search that found it.

11. Sensitivity to priors, kernels and encodings

Bayesian optimisation results should be challenged against reasonable alternative modelling choices. A candidate that is optimal only under one arbitrary categorical encoding or one strongly informative length-scale prior is not yet robust. Sensitivity analysis should include weaker and stronger priors, alternative kernels for mixed variables, and removal of historical data that may carry domain mismatch. The resulting change in candidate ranking is part of the uncertainty of the decision. Categorical chemistry choices deserve particular attention because numerical encodings can impose artificial distance. Descriptor-based representations may be useful when descriptors are physically meaningful and measured consistently, but a descriptor itself can introduce a hidden model of similarity. The study should therefore report both the representation and the evidence supporting its use.

12. Sequential design versus qualification design

The experiment design that is optimal for learning a response surface is not necessarily the design that is strongest for qualification. Sequential selection concentrates observations in informative regions and can leave sparse coverage elsewhere. Once a candidate is selected, confirmation should switch to a predefined design focused on repeatability, sensitivity to critical process variables and acceptance limits rather than continued optimisation. This separation also protects against winner selection bias. The best observed candidate in an adaptive campaign has benefited from repeated comparison against many alternatives; its measured value can therefore be optimistic. Fresh independent replicates provide a cleaner estimate of its actual performance.

13. Failure signatures and diagnostics

Warning signs include repeated acquisition of nearly identical points without reduction in predictive uncertainty, strong disagreement between cross-validation and prospective prediction, feasibility failures concentrated in one categorical branch, and abrupt optimum movement after a single observation. These patterns may indicate misspecified covariance structure, unmodelled regime boundaries or measurement drift. Another failure mode is apparent efficiency created by excluding unsuccessful experiments from the denominator. Experimental efficiency should count preparation failures, non-cures, out-of-range measurements and other executed trials unless they are proven executor faults and separately reported.

14. Evidence gaps and research priorities

Photopolymer Bayesian optimisation needs benchmark datasets in which formulation and process variables are both represented and where failed experiments remain visible. Most generic materials benchmarks do not include cure-history variables, post-processing state or optical-process interactions that are central to photopolymerisation. Future studies should compare constrained BO against simpler sequential designs under identical experimental budgets and include an independent holdout campaign. Demonstrating that a sophisticated optimiser beats random or factorial search only in retrospective simulation is weaker evidence than prospective performance with a real executor and real measurement uncertainty.

References

1. Jin Y, Kumar PV. Bayesian optimisation for efficient material discovery: a mini review. *Nanoscale*. 2023;15:10975-10984. doi:10.1039/D2NR07147A.
2. Kusne AG, Yu H, Wu C, Zhang H, Hatrick-Simpers J, DeCost B, Sarker S, Oses C, Toher C, Curtarolo S, Davydov AV, Agarwal R, Bendersky LA, Li M, Mehta A, et al. On-the-fly closed-loop materials discovery via Bayesian active learning. *Nature Communications*. 2020;11:5966. doi:10.1038/s41467-020-19597-w.
3. Lookman T, Balachandran PV, Xue D, et al. Active learning in materials science with emphasis on adaptive sampling using uncertainties for targeted design. *npj Computational Materials*. 2019;5:21. doi:10.1038/s41524-019-0153-8.
4. Fadhel A, Zuckschwerdt NW, Deshwal A, Bose S, Bandyopadhyay A, Doppa J. Discovery of Feasible 3D Printing Configurations for Metal Alloys via AI-Driven Adaptive Experimental Design. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2026;40(47):39939-39948. doi:10.1609/aaai.v40i47.41428.

Competing interests

The author is Founder, CEO and CTO of 3Dresyns by Resyner Technologies S.L.; ER-402 contains no comparative recommendation between commercial products.

Data availability

ER-402 is a literature-based engineering review; no new experimental dataset was generated.