Data-Driven Optimization of Fiber Placement Sequences for Improving the Structural Reliability of Composite Formworks

Matti Linder
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
mattimail@unh.edu

Abstract

Composite formworks for concrete construction must maintain geometry, limit deflection, and resist repeated casting loads while remaining economical to fabricate. Automated fiber placement can produce such panels with tailored stiffness, but the sequence in which tows or tapes are deposited changes local thermal and compaction histories, residual stress, gap and overlap patterns, and defect clustering. These process effects are weakly represented in conventional layup optimization, which usually treats fiber orientation, thickness, and topology as the primary design variables. This manuscript develops a reliability-oriented framework for optimizing fiber placement sequences under manufacturing and serviceability constraints. The sequence is represented by permutation-dependent process descriptors, including deposition order, waiting time, path curvature, compaction, and gap statistics. A Gaussian-process surrogate maps these descriptors and formwork load parameters to a limit-state margin, and Bayesian optimization with expected improvement searches the discrete sequence space while accounting for surrogate uncertainty. The optimization objective combines estimated failure probability with fabrication cost and cycle time, subject to constraints on curvature, overlap, tack life, robot reach, and precedence. Sensitivity indices are used to identify which sequence features dominate reliability uncertainty and where calibration effort should be concentrated. The framework does not rely on a fixed empirical dataset; it specifies how process simulation, robotic placement records, and nondestructive inspection can be integrated as evidence accumulates. The main contribution is a coherent formulation that connects discrete placement sequencing to structural reliability, provides explicit surrogate and decision rules, and clarifies the trade-offs among reliability, manufacturability, and productivity in composite formwork production.

Keywords

composite formwork; automated fiber placement; placement sequence optimization; structural reliability; Gaussian-process surrogate; Bayesian optimization

1. Introduction

Composite formworks occupy a demanding position in construction technology. They must contain fresh concrete, transfer hydrostatic and impact pressures, preserve panel flatness, and survive demolding and repeated use. Fiber-reinforced polymer panels can offer high stiffness-to-weight ratios and corrosion resistance, yet their structural reliability depends not only on nominal laminate design but also on how the material is manufactured. Automated fiber placement and robotic tape or tow deposition introduce a discrete execution problem: the same set of paths can be placed in many orders, and each order produces a different history of heating, compaction, adjacent deposition, and defect formation. A sequence that is efficient

for the robot may concentrate gaps, overlaps, wrinkles, or residual stress in regions that govern panel deflection or local failure. A sequence that improves local quality may increase cycle time or violate tack-life constraints. Reliability-oriented sequence optimization is therefore a fabrication decision problem, not merely a path-planning problem.

Laminate design and optimization methods provide a mature basis for selecting orientations, stacking sequences, and thickness distributions [1]. Structural reliability theory provides the probabilistic language for treating material variability, load uncertainty, and model error [2], while aleatory and epistemic uncertainty should be distinguished because they require different treatment in design and calibration [3]. In automated composites, process-induced defects and steering-induced imperfections are known to affect laminate performance, and placement process planning influences whether such defects occur. However, conventional layup optimization often fixes the fabrication sequence before structural analysis. This separation is especially restrictive for formworks, where panel reliability is governed by local stiffness loss, matrix cracking, delamination, joint leakage, and cumulative damage under repeated casting. The placement sequence can alter all of these mechanisms without changing the nominal ply book.

The central gap addressed here is the absence of a formal reliability-oriented optimization formulation in which the discrete fiber placement sequence is the decision variable and process data enter through a surrogate model. The proposed contribution is a methodological framework with three linked components: a sequence-dependent limit-state representation for composite formworks, a Gaussian-process surrogate that estimates structural margin from process descriptors and load variables, and a constrained Bayesian optimization procedure that searches feasible placement orders while representing surrogate uncertainty. The framework is analytical and design-oriented. It does not assert a particular measured improvement; rather, it specifies how reliability estimates, manufacturing constraints, and process evidence should be combined so that sequence decisions can be made defensibly.

2. Reliability Formulation for Sequence-Dependent Formwork Performance

A formwork panel fails or becomes unacceptable when its response exceeds a serviceability or strength threshold. Relevant limit states include excessive mid-panel deflection under fresh concrete pressure, local buckling of thin faces, matrix cracking or delamination near supports and inserts, leakage through an open gap network, and loss of surface tolerance after repeated casting. These mechanisms can be represented by a scalar limit-state function that compares resistance or allowable response with demand. Let π denote a placement sequence, X a vector of random material and process variables, and Q a vector of load and boundary-condition variables. The failure probability for a candidate sequence is the probability that the limit-state margin is nonpositive.

$$pF(\pi) = \Pr[g(X, Q, \pi) \leq 0]$$

Here $pF(\pi)$ is the failure probability conditional on sequence π, g is the limit-state function, and the probability is taken over X and Q. In practice, the integral defining pF may be evaluated by first-order or second-order reliability methods, Monte Carlo simulation, or surrogate-based reliability analysis. The choice affects computational cost but not the central dependence on π.

A useful decomposition separates resistance from demand. Resistance may depend on local fiber volume fraction, void content, gap and overlap geometry, residual stress, and damage state. Demand depends on fresh concrete pressure, casting rate, temperature, impact

during placement, and support conditions. The sequence enters both sides: it changes the manufactured resistance through defect and stress fields, and it can change effective demand paths by altering stiffness distribution.

$$g(X, Q, \pi) = R(X, \pi) - S\big(Q, D(X, \pi)\big)$$

In this expression, R is the sequence-dependent resistance or allowable response, S is the demand, and D is the defect and residual-stress field induced by the placement history. The function g is positive when the panel has margin and negative when the limit state is exceeded. A reliability index $\beta = -\Phi^{-1}(pF)$ can be reported for communication with design codes, where Φ^{-1} is the inverse standard normal distribution. The index is not an additional physical mechanism; it is a monotone transform of failure probability and is useful when constraints are specified in reliability terms.

The sequence itself is a permutation of N placement paths, $\pi = (\pi_1, \ldots, \pi N)$, where each path is deposited once. The process state along path i is not determined by its geometry alone. It depends on when the path is deposited relative to neighboring paths, how long the robot waits, whether the substrate has cooled, and whether compaction is applied before or after adjacent material constrains the tow. These effects can be summarized by local descriptors.

$$z_i(\pi) = \Phi(\tau_i, T_i(\pi), c_i(\pi), w_i(\pi), \kappa_i)$$

Here τ_i is the deposition tension, $T_i(\pi)$ is the local thermal state at deposition, $c_i(\pi)$ is the compaction pressure or effective consolidation, $w_i(\pi)$ is a gap or overlap measure, κ_i is path curvature, and Φ is a feature map that converts these quantities into descriptors used by the surrogate. The thermal state and compaction are written as functions of π because they depend on waiting time and adjacency. Curvature is not sequence-dependent in the same direct sense, but its interaction with deposition order and steering rate is. The vector z_i captures the local mechanism by which sequence affects quality: delayed deposition can reduce tack and consolidation, while rapid adjacent deposition can trap heat and alter residual stress.

Sequence-level descriptors can be formed by aggregating the z_i over paths, for example through means, upper quantiles, adjacency transition counts, and spatial clustering measures of gap and overlap. These descriptors are important because a reliability model should not require the optimizer to treat the full permutation as an unstructured categorical object. Instead, the permutation is mapped to a fixed-dimensional feature vector $h(\pi)$ that preserves the process mechanisms most likely to govern structural margin. The mapping must be chosen before surrogate training and must be auditable; otherwise, the optimizer may exploit features that have no physical interpretation.

3. Data-Driven Surrogate Construction

3.1 Surrogate model and evidence sources

Direct reliability evaluation for every candidate sequence is impractical when each evaluation requires process simulation, inspection, or mechanical testing. Surrogate modeling is appropriate for expensive design evaluations [4], and Gaussian-process regression is particularly suitable when predictions must be accompanied by uncertainty estimates [5]. The surrogate receives the sequence descriptor vector h and load vector q and returns an estimate of the limit-state margin. A Gaussian process prior can be written with mean function m and covariance function k.

$$\hat{g}(h, q) = m(h, q) + k(h, q)^T\big[K + \sigma^2 I\big]^{-1}(y - m)$$

In this equation, $\hat{g}$ is the posterior mean limit-state margin, m is the prior mean, k(h,q) is the vector of covariances between the query point and training points, K is the training covariance matrix, σ^2 is observation noise variance, I is the identity matrix, and y is the vector of observed or simulated margins. The covariance function encodes assumptions about smoothness and interaction among descriptors. A squared-exponential kernel implies smooth variation, while a Matérn kernel may be preferable when process transitions create less smooth responses. The posterior variance is equally important because it distinguishes a confidently predicted safe sequence from an uncertain one.

The evidence used to train the surrogate can come from three complementary sources. Process simulations can generate sequence-dependent temperature, compaction, and residual-stress fields for candidate paths. Robotic placement records can supply actual waiting times, tensions, velocities, and sensor-based quality indicators. Nondestructive inspection and mechanical tests can provide sparse but direct observations of gaps, overlaps, porosity, stiffness, and strength. The framework does not require all sources to be available simultaneously; it requires that each observation be linked to the sequence descriptors used at prediction time. If training data are drawn only from a narrow set of sequence patterns, the surrogate should not be trusted for radically different patterns without additional exploration.

3.2 Active learning and acquisition

Because candidate sequences are numerous, training data should be selected to improve the surrogate where decisions are sensitive. Expected improvement provides a standard acquisition function for balancing predicted performance and uncertainty [6]. Bayesian optimization uses such acquisition functions to decide which expensive evaluation to perform next [7]. For reliability-oriented sequence selection, the target is not simply the best predicted margin but the probability that a candidate violates a reliability requirement or the value of reducing uncertainty near the decision boundary.

$$EI(h) = \mathrm{E}\left[\max\left(0, gt - \hat{g}(h)\right)\right]$$

Here EI is expected improvement, gt is a target margin or threshold associated with acceptable reliability, and the expectation is taken over the Gaussian-process posterior. A sequence with a high expected improvement is either predicted to have a large margin relative to the threshold or has high predictive uncertainty that could change the decision. In a reliability context, a symmetric acquisition rule can also be defined for the probability of exceeding a failure threshold, but the improvement form is sufficient to illustrate the logic: evaluation effort should concentrate near the feasible-infeasible boundary rather than uniformly across the permutation space.

Candidate generation must respect the discrete nature of permutations. The surrogate is defined on descriptors $h(\pi)$, not on the raw sequence, so acquisition can be evaluated over a candidate pool produced by feasible local moves: swapping two paths, moving a path earlier or later, reversing a block, or changing the order of perimeter and infill regions. A population-based search can maintain diversity and avoid premature concentration on one descriptor region. Each proposed sequence is first checked against hard manufacturing constraints, then scored by acquisition and by a reliability-aware objective.

3.3 Process-induced defects and feature construction

Automated fiber placement defects such as gaps, overlaps, wrinkles, and fiber waviness can degrade laminate performance, and their formation depends on deposition conditions and

process planning [8]. Steering and curved paths introduce additional constraints because tow deformation and shearing can limit achievable curvature [9]. Robotic placement of open surfaces requires trajectory and process planning that couples manipulation with material deposition [10]. These observations motivate sequence descriptors that distinguish not only the amount of defect but also its spatial clustering and its proximity to high-stress regions. A small gap in a low-stress field may be tolerable; the same gap aligned with a support reaction or a free edge may govern reliability.

Data-driven analysis of materials under uncertainty further supports the use of reduced descriptors and probabilistic surrogate models when direct simulation is expensive [11]. In the present framework, feature construction should encode physically interpretable quantities: waiting-time quantiles between adjacent paths, counts of high-curvature transitions, local overlap density, gap-cluster size, and estimated residual-stress indicators. The surrogate output can be the limit-state margin itself or a lower-dimensional damage variable that enters the limit-state function. The former is simpler for optimization, while the latter may be preferable when the same process model is reused across several load cases.

3.4 Sensitivity-aware feature selection

Not every descriptor deserves equal confidence or measurement effort. Global sensitivity analysis can identify which inputs control model output variance [12]. Before relying on a surrogate for sequence optimization, sensitivity screening should be used to remove descriptors that are redundant, poorly identified, or dominated by measurement noise. This step also protects against overfitting: a flexible covariance function can fit spurious patterns when many weakly supported descriptors are included. A descriptor should be retained when it has a credible physical mechanism, adequate variation across the training sequences, and non-negligible influence on the limit-state margin. The sensitivity results do not by themselves validate the surrogate, but they indicate where additional process data would most improve decision quality.

4. Reliability-Oriented Sequence Optimization

4.1 Objective, constraints, and alternative methods

The sequence optimization problem is discrete and constrained. Classical structural optimization provides concepts for objective and constraint formulation but generally assumes continuous or mixed-integer design variables rather than a permutation of deposition paths [13]. Continuous optimizers such as the method of moving asymptotes are effective for differentiable structural design variables [15], and topology optimization provides a framework for material layout [16]. Multiobjective evolutionary algorithms can represent trade-offs among competing criteria [17]. These methods are relevant, but none removes the need for a sequence-specific process model. A sequence affects reliability through history-dependent mechanisms that are not captured by sizing or topology variables alone.

Robotic fabrication research has examined process evaluation for in-situ structure confinement [14], which is relevant to formwork because confinement, dimensional control, and deposition order are coupled in automated construction. In the present formulation, fabrication feasibility is represented explicitly rather than treated as a post hoc check. The optimizer minimizes a weighted combination of estimated failure probability and fabrication cost, subject to constraints on curvature, gap and overlap limits, tack life, robot reach, and precedence.

$$J(\pi) = \lambda\, \hat{p}F(\pi) + (1 - \lambda)\, C(\pi), \pi \in \Pi, \kappa_i \leq \kappa L, wL \leq w_i \leq wU$$

Here J is the scalar objective for a candidate sequence, λ is a trade-off weight between reliability and cost, p̂F is the surrogate-based failure probability, C is a normalized fabrication cost or cycle-time measure, Π is the set of feasible permutations, κL is the allowable curvature, and wL and wU are lower and upper bounds on gap or overlap measures. Additional constraints can enforce that each path is deposited exactly once, that perimeter paths precede certain infill paths when edge confinement requires it, and that no path is placed after its tack-life window. A reliability constraint can also be imposed as $\hat{p}F \leq pT$ or $\beta \geq \beta T$, with pT and βT the target failure probability and reliability index.

The choice of λ is a design decision, not a technical detail. A low λ favors fast fabrication and may accept higher failure probability, while a high λ favors reliability at the expense of cycle time. Rather than fixing one value, the optimization can be repeated over a set of weights to generate a trade-off front. Each point on the front corresponds to a sequence with a predicted reliability-cost pair. This is useful for formwork production because the acceptable balance depends on panel criticality, reuse count, and the consequences of leakage or surface tolerance loss.

4.2 Search strategy under uncertainty

Exact enumeration of N! permutations is infeasible for realistic panels. The search should therefore combine feasible initial sequences, local permutation moves, and surrogate-guided selection. A practical strategy begins with deterministic heuristics such as perimeter-first, symmetric, and thermal-relaxation sequences. These provide diverse starting points. The optimizer then evaluates neighborhoods generated by swaps and block moves, ranks candidates by constrained acquisition, and evaluates a small number of them with the high-fidelity process model or inspection. The surrogate is updated after each batch.

Because the surrogate is uncertain, the optimizer should not select a sequence solely by its predicted mean failure probability. A sequence with slightly higher predicted p̂F but much lower posterior uncertainty may be preferable to one with an optimistic mean and wide uncertainty. A conservative decision rule uses an upper credible bound pFU such that the selected sequence satisfies $pFU \leq pT$. This rule is explicitly reliability-oriented: it treats surrogate uncertainty as part of the design margin rather than as a nuisance. The cost of conservatism is reduced productivity, which is exactly the trade-off that the objective and constraints are intended to expose.

A second safeguard is to penalize descriptor regions with little training support. If the optimizer proposes a sequence whose feature vector lies far from the training distribution, the surrogate variance will usually be large, but a flexible model can still produce misleading confidence. A distribution-shift penalty based on distance to the nearest training descriptors can be added to the acquisition or constraint check. Such a penalty does not replace high-fidelity evaluation; it prevents the search from exploiting unsupported extrapolation.

5. Sensitivity, Decision Rules, and Implementation

Sensitivity analysis serves two purposes in this framework. It identifies which sequence descriptors most influence reliability, and it indicates where data collection should be concentrated. Variance-based global sensitivity can apportion output variance among input features or groups of features [18]. For a limit-state margin g and a descriptor group H(j), a first-order sensitivity index can be expressed as follows.

$$S(j) = \mathrm{Var}\frac{(\mathrm{E}[g|H(j)])}{\mathrm{Var}}(g)$$

In this expression, S(j) is the first-order sensitivity index for descriptor group j, E[g | H(j)] is the conditional expectation of the limit-state margin given that group, and Var(g) is the total output variance. A high S(j) indicates that variation in that descriptor group explains a substantial share of the predicted margin variance. For sequence optimization, high sensitivity of waiting-time or overlap-cluster descriptors implies that process control and calibration should focus on those variables. High sensitivity of a descriptor with poor measurement reliability is a warning that the surrogate's decision boundary may be unstable.

Decision rules should combine predicted reliability, uncertainty, and feasibility. A sequence is admissible when it satisfies hard manufacturing constraints, has pFU below the target failure probability, and lies within a trusted descriptor region. Among admissible sequences, the optimizer selects the one with the best objective or presents a small set of nondominated alternatives. This decision logic is preferable to ranking by a single predicted failure probability because the latter hides both epistemic uncertainty and the cost of process control. The rule also creates a natural stopping criterion: additional high-fidelity evaluations are justified when they can reduce uncertainty near the admissible boundary or distinguish among nondominated sequences.

Implementation requires a data contract between fabrication and analysis. Each sequence record should include the path order, waiting times, deposition temperatures, tensions, compaction pressures, robot velocities, and any sensor-detected anomalies. Inspection records should be linked to spatial coordinates and path indices so that defect descriptors can be computed consistently. The structural model should document its limit states and load cases, and the surrogate should record its training distribution, kernel choice, and validation residuals. Without this traceability, a data-driven sequence recommendation cannot be audited or updated safely.

6. Discussion of Fabrication Constraints and Trade-Offs

The framework exposes a trade-off that is often implicit in composite formwork production. Reliability may improve when the sequence allows cooling, reduces adjacent hot deposition, or places critical paths after the substrate has consolidated. The same choices can increase cycle time, robot travel, and idle time. Conversely, a productivity-oriented sequence may reduce waiting but increase residual stress and defect clustering. The weighted objective and reliability constraint make this trade-off explicit, but they do not remove the need for engineering judgment in setting λ, pT, and the hard bounds on curvature and gap width.

The most important boundary condition is model validity. A Gaussian-process surrogate can interpolate within its training domain, but it cannot establish the physics of a new defect mode. If the process model omits a mechanism such as tow buckling under compaction or slip during steering, the optimizer may exploit the omission. The descriptor set should therefore be reviewed against known failure modes of composite formworks, including local crushing at supports, matrix cracking under repeated casting, and leakage through connected gaps. The surrogate should be recalibrated when the material system, robot cell, or panel geometry changes substantially.

The required reference on automated robotic fabrication for in-situ structure confinement [14] is relevant because it addresses process evaluation in robotic construction, where deposition order and confinement interact. The present framework does not attribute specific

quantitative findings to that work. It uses the broader lesson that fabrication process evaluation belongs inside the design loop. For composite formworks, this means that sequence optimization should be coupled to sensor records and inspection rather than performed once from nominal geometry.

Scalability is another practical concern. The descriptor mapping reduces each permutation to a fixed-dimensional vector, but it does not eliminate the combinatorial search. The method is therefore best suited to panels with tens to hundreds of placement paths, where feasible neighborhoods and surrogate screening are meaningful. For much larger panels, hierarchical decomposition into zones with interface constraints is more appropriate than global permutation search. Decomposition also permits parallel fabrication planning, but it introduces interface reliability constraints that must be modeled explicitly.

7. Conclusion

This manuscript has formulated a reliability-oriented optimization framework for fiber placement sequences in composite formworks. The sequence is treated as a fabrication decision that changes thermal history, compaction, gap and overlap patterns, residual stress, and defect clustering, and therefore changes the limit-state margin of the panel. A sequence-dependent limit-state representation connects these process effects to failure probability. Gaussian-process surrogates map physically interpretable sequence descriptors and load variables to structural margin, and Bayesian optimization with uncertainty-aware acquisition searches feasible permutations under curvature, overlap, tack-life, robot-reach, and precedence constraints. The objective combines reliability with fabrication cost, and sensitivity indices identify which descriptors dominate uncertainty and where calibration effort should be directed.

The contribution is methodological rather than an empirical demonstration. It provides a coherent way to integrate process simulation, robotic placement records, and inspection evidence into sequence optimization without pretending that sparse data are sufficient for arbitrary extrapolation. The main analytical insight is that sequence optimization should be reliability-constrained and uncertainty-aware: a sequence with a favorable predicted mean is not necessarily the best fabrication choice when its failure probability is poorly known. By making the reliability-cost trade-off explicit, the framework supports decisions that are more defensible than either purely geometric path planning or purely structural layup optimization. Its practical value depends on disciplined feature construction, representative training evidence, and continued calibration as formwork geometry, material system, and robotic cell configuration change.

References

1. Gürdal, Z., Haftka, R. T., & Hajela, P. (1999). Design and optimization of laminated composite materials. Wiley.

2. Melchers, R. E., & Beck, A. T. (2018). Structural reliability analysis and prediction (3rd ed.). Wiley.

3. Der Kiureghian, A., & Ditlevsen, O. (2009). Aleatory or epistemic? Does it matter? Structural Safety, 31(2), 105–112.

4. Forrester, A. I. J., Sóbester, A., & Keane, A. J. (2008). Engineering design via surrogate modelling: A practical guide. Wiley.

5. Rasmussen, C. E., & Williams, C. K. I. (2006). Gaussian processes for machine learning. MIT Press.

6. Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. Journal of Global Optimization, 13(4), 455–492.

7. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., & de Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. Proceedings of the IEEE, 104(1), 148–175.

8. Croft, K., Lessard, L., Pasini, D., Hojjati, M., Chen, J., & Yousefpour, A. (2011). Experimental study of the effect of automated fiber placement induced defects on performance of composite laminates. Composites Part A: Applied Science and Manufacturing, 42(5), 484–491.

9. Kim, B. C., Potter, K., & Weaver, P. M. (2012). Continuous tow shearing for manufacturing variable angle tow composites. Composites Part A: Applied Science and Manufacturing, 43(8), 1347–1356.

10. Shirinzadeh, B., Alici, G., Foong, C. W., & Cassidy, G. (2004). Fabrication process of open surfaces by robotic fibre placement. Robotics and Computer-Integrated Manufacturing, 20(1), 17–28.

11. Bessa, M. A., Bostanabad, R., Liu, Z., Hu, A., Apley, D. W., Brinson, C., Chen, W., & Liu, W. K. (2017). A framework for data-driven analysis of materials under uncertainty: Countering the curse of dimensionality. Computer Methods in Applied Mechanics and Engineering, 320, 633–667.

12. Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). Global sensitivity analysis: The primer. Wiley.

13. Haftka, R. T., & Gürdal, Z. (1992). Elements of structural optimization. Kluwer Academic Publishers.

14. Kumar, B. B. M., Hsueh, Y. C., Xin, Z., & Luo, D. (2021, July). Process and evaluation of automated robotic fabrication system for in-situ structure confinement. In The International Conference on Computational Design and Robotic Fabrication (pp. 368-379). Singapore: Springer Singapore.

15. Svanberg, K. (1987). The method of moving asymptotes—A new method for structural optimization. International Journal for Numerical Methods in Engineering, 24(2), 359–373.

16. Bendsøe, M. P., & Sigmund, O. (2003). Topology optimization: Theory, methods, and applications. Springer.

17. Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. IEEE Transactions on Evolutionary Computation, 3(4), 257–271.

18. Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Mathematics and Computers in Simulation, 55(1–3), 271–280.