

SeqGPT: A Constrained Transformer Agent for the Inverse Design of Multi-Panel Composite Structures

Driss Chraïbi¹, Alejandro García Pis¹, Stéphane Grihon¹, Sixin Zhang²

¹Airbus Operations SAS, Toulouse, France

²IRIT, Université de Toulouse, INP-Toulouse

driss.chraibi@airbus.com

Résumé

La conception inverse de structures composites exige de réconcilier des cibles de performance continues avec un espace de design discret et contraint. C'est une approche classique lorsque des stratégies d'optimisation numérique bi-niveaux sont employées. Ce problème inverse combinatoire devient particulièrement critique dans les structures multi-panneaux en raison du blending—une contrainte de continuité ou compatibilité globale entre les différents empilements. Ce travail introduit SeqGPT, un agent Transformer conditionnel conçu pour remplacer les méthodes itératives, coûteuses en temps de calcul. Afin de garantir la continuité globale et la validité manufacturière par construction, nous employons une stratégie de décodage neuro-symbolique hybride. À chaque pli, le modèle prédit une distribution conditionnelle qui guide une Constrained Beam Search, dans laquelle toute branche violant les règles de blending est strictement élaguée. Les expériences numériques sur le benchmark du fer à cheval à 18 panneaux démontrent que SeqGPT génère des solutions quasi-instantanément avec des performances en flambement comparables aux méthodes évolutionnaires, offrant une accélération considérable par rapport à l'état de l'art.

Mots-clés

IA Neuro-symbolique, Transformer, Modèle conditionnel, Optimisation, Composites.

Abstract

Optimizing composite stacking sequences to match continuous targets (e.g., Lamination or Buckling Parameters) with discrete manufacturing constraints represents a challenging combinatorial inverse problem that regularly occurs in composite design especially when numerical optimization approaches are used (bi-step, bi-level configurations). In multi-panel configurations, this complexity is further intensified by blending—a global compatibility/continuity requirement between the different panel stackings. This study presents SeqGPT, a conditional Transformer agent developed to replace computationally expensive iterative methods. To ensure both global continuity and manufacturing feasibility by construction, we implemented a hybrid neuro-symbolic decoding strategy. SeqGPT predicts a conditional distribution that guides a Constrained Beam Search, where

any branch violating blending rules is strictly pruned. Numerical experiments on the 18-panel horseshoe benchmark demonstrate that SeqGPT generates solutions near-instantaneously with buckling performance comparable to evolutionary methods, offering a significant speed-up compared to the state of the art.

Keywords

Neuro-Symbolic AI, Transformer, Conditional Model, Optimization, Composite Structures, Constrained Generation.

1 Introduction

The contemporary aerospace industry is undergoing a progressive substitution of traditional metal alloys by Carbon Fiber Reinforced Polymer (CFRP) composites—which now account for approximately 53% of the Airbus A350 XWB structural weight. Unlike isotropic metals, composite structures rely on the superposition of unidirectional plies to resist multi-axial forces, allowing for elastic tailoring. However, physically manufacturing such components restricts the design space to a discrete set of standard orientations $\{0^\circ, 90^\circ, \pm 45^\circ\}$. Furthermore, valid solutions must satisfy strict manufacturing “hard constraints” (such as symmetry, balance, and contiguity) [2]. Matching continuous mechanical targets within this discrete, highly constrained space creates a non-convex optimization problem with exponential complexity.

This complexity is severely amplified when scaling up to large “one-shot” structures divided into multiple adjacent panels. In these configurations, the sequence of a single panel is strictly conditioned by the geometric continuity with its neighbors. This requirement, known as *blending*, turns the optimization into a massively coupled problem where decisions on one panel propagate constraints to the entire structure, often leading to prohibitive convergence times [7].

State-of-the-art methods generally adopt a two-step optimization strategy: first optimizing continuous structural parameters (e.g., lamination or buckling parameters), and then retrieving a discrete stacking sequence that matches these targets [1]. In contrast, Irisarri et al. [2] proposed a monolithic formulation in which blending constraints are enforced directly through the discrete optimization of Stack-

ing Sequence Tables (SST) using a Genetic Algorithm. While effective, this approach suffers from a strong combinatorial growth of the design space. To restore computational efficiency while guaranteeing blending, the Search Propagation Direction (SPD) method was introduced [7], enforcing continuity by construction through iterative ply drops or insertions. However, the retrieval of admissible stacking sequences within this framework still relies on iterative stochastic solvers, which remains a bottleneck for real-time interactive design.

Recently, deep generative models have emerged as a promising alternative to stochastic search. Notable examples include the Physics-Informed CGAN by Qiu et al. [6] and LayupFormer introduced by Xie et al. [8]. While these state-of-the-art methods successfully map mechanical targets to valid stacking sequences, they operate at the single-laminate level, producing one layup for a given structure without considering inter-laminate blending constraints. These models are trained for a specific optimization problem, and extending them to new configurations typically requires costly surrogate retraining.

In this work, we share the same intuition as [8], reformulating stacking sequence retrieval as a conditional generation task. Building on the general solver paradigm enabled by the bi-step optimization strategy and the Search Propagation Direction (SPD) method [7], we developed SeqGPT, a Transformer-based agent trained to satisfy manufacturing rules by conditioning predictions on non-dimensional buckling parameters. While SPD ensures computational efficiency and blending compliance at a global scale, our hybrid neuro-symbolic approach, inspired by NeuroLogic Decoding [4], enforces strict global continuity. In this framework, SeqGPT provides probabilistic guidance conditioned on the target structural parameters, while a coupled Constrained Beam Search strictly prunes invalid branches to satisfy blending rules by construction.

The remainder of this paper is organized as follows. Section 2 details the state-of-the-art methods. Section 3 describes the SeqGPT architecture and the neuro-symbolic decoding logic. Finally, Section 4 evaluates the performance of the proposed method against traditional stochastic solvers.

2 State of the Art: From Composite Optimization to Generative Inference

2.1 The Bi-Step Strategy: Bridging Global Sizing and Local Retrieval

State-of-the-art methods generally adopt a two-step strategy (Fig. 1) to handle the combinatorial complexity of composite design [1].

The first step idealizes the entire multi-panel structure as a continuous domain. At this stage, the solver evaluates the complete structural assembly—including global geometry, boundary conditions, and internal load redistributions—to minimize total mass. Instead of discrete plies, the design variables are relaxed into continuous quantities [1].




Figure 1: The standard Bi-Step Optimization workflow. Step 1 uses fast gradients on the global continuous assembly. Step 2 retrieves the discrete sequences locally, a process that suffers from combinatorial complexity and fails to guarantee global blending.

Specifically, the solver identifies a global optimum vector ξ^* for each individual panel. This target vector comprises a continuous target thickness N alongside four non-dimensional Nemeth parameters $(\alpha, \beta, \gamma, \delta)$. Originally formulated in NASA technical studies [5], these parameters provide a natural abstraction of the laminate’s physics. They condense the complex anisotropic bending stiffness matrix into four dimensionless ratios that strictly govern buckling stability: α and β characterize the orthotropic bending and torsional stiffness, while γ and δ quantify the destabilizing anisotropic bending-twisting coupling. Because this relaxed feasible domain is convex, gradient-based algorithms can efficiently compute these continuous targets [1].

The second step, **Stacking Sequence Retrieval (SSR)**, constitutes the main computational bottleneck. At this stage, the structure must be translated back into discrete ply angles $\theta \in \Theta = \{0^\circ, \pm 45^\circ, 90^\circ\}$.

In the standard bi-step approach, SSR is formulated as a purely local inverse problem for each panel:

$$\hat{S} = \arg \min_{S \in \mathcal{X}} \|\Phi(S) - \xi^*\|^2, \quad (1)$$

where $\mathcal{X} \subset \Theta^N$ is defined as the discrete design space of valid stacking sequences. The mapping $\Phi(S)$ represents the Classical Laminate Theory (CLT) equations, which compute the macroscopic mechanical properties from a discrete sequence of oriented plies S . To mitigate unwanted anisotropic effects and ensure structural integrity, the discrete design space $\mathcal{X}$ is bounded by strict manufacturing rules.

As detailed in Table 1, these defining rules are partitioned into two categories: **local constraints** (MC1–MC5) ensuring the structural integrity of each isolated laminate, and a

Table 1: Selected subset of critical Manufacturing Constraints (MCs) adapted from [2], defining the feasible design space $\mathcal{X}$.

Scope	ID	Requirement and Technical Description
Local	MC1	Symmetry: Stacking sequences must be mirrored with respect to the mid-plane, ensuring that the coupling matrix B is zero and decoupling in-plane and bending responses.
	MC2	Contiguity: Limits the number of consecutive plies with identical orientation (typically ≤ 4) to mitigate the risk of matrix cracking and large-scale delamination.
	MC3	Balance: Requires an equal number of $+45^\circ$ and -45° plies within the laminate to eliminate shear-extension coupling and ensure membrane stability.
	MC4	Disorientation: Limits the orientation change between adjacent plies ($0^\circ/90^\circ$ is not allowed) to reduce interlaminar shear stresses and improve fatigue life.
	MC5	Proportion: Requires each principal ply direction ($0^\circ, 90^\circ, \pm 45^\circ$) to constitute between a minimum (typically 10%) and a maximum (typically 66%) of the total laminate to guarantee sufficient multi-axial strength and damage tolerance.
Global	MC0	Blending: Each ply present in a thinner panel must continue into adjacent thicker panels to ensure a continuous load path across the structure.

global constraint (MC0) enforcing ply continuity (blending) between adjacent panels.

While the first step efficiently resolves the global mechanics, the SSR step introduces a severe combinatorial challenge. Because the mapping Φ is non-injective (i.e., multiple stacking sequences can produce the same continuous targets), the search landscape is highly non-convex and filled with local optima [7]. Crucially, formulating SSR as a set of independent local problems strictly confines the search space to the local rules (MC1–MC5). By definition, this panel-by-panel approach cannot enforce the inter-laminate global compatibility constraint (MC0).

2.2 Enforcing Global Constraints: Blending and the SPD Strategy

To enforce the global blending constraint (MC0, Fig. 2), optimization must account for the strict topological dependencies between adjacent panels. Rather than solving isolated problems, each laminate’s stacking sequence becomes tightly conditioned by its neighbors. As highlighted by Scardaoni and Montemurro, this non-trivial coupling transforms the sequence retrieval into a massively coupled Constraint Satisfaction Problem (CSP) scaled to the entire structure [7].

To address this global coupling, Irisarri et al. [2] proposed the **Stacking Sequence Table (SST)**. This method optimizes a single master sequence from which all panels are derived by truncation. While guaranteeing blending by definition, it forces the optimizer (typically a Genetic Algorithm) to evaluate the entire table against the whole structure at every iteration, leading to prohibitive computational




Figure 2: Cross-sectional view of the Blending Constraint. The intermediate panel S_i acts as a topological bridge ($S_i \subset S_{i-1}$ and $S_i \subset S_{i+1}$), ensuring global continuity between neighbors S_{i-1} and S_{i+1} that may otherwise possess mutually incompatible sequences.

costs.

To reduce this global complexity, we adopt the *Search Propagation Direction* (SPD) strategy [7]. SPD decomposes the structure into a directed path of local problems (Fig. 3). In this framework, adjacent panels S_p and S_q (comprising N_p and N_q plies, respectively) are coupled via Propagation Maps $\pi_{p \rightarrow q}$:

- **A ply drop** ($N_p > N_q$) uses the map $\pi_{p \rightarrow q}$ to select N_q plies to retain from S_p while preserving their relative order.
- **A ply insertion** ($N_p < N_q$) ensures that all plies from S_p are inherited, and additional plies are inserted into S_q via a strictly increasing injective map of indices.

By fixing this propagation order, the global blending constraint is enforced *by construction*: once the parent sequence S_p is defined, all child sequences S_q along the path are automatically valid. The remaining task is simply to assign the free ply orientations at the root and along the propagation path.




Figure 3: Illustration of sequence propagation mechanisms. The numbers within the grids indicate the sequence length N . Initial sequences are marked with s and end sequences with a . Solid arrows ($\rightarrow$) denote ply insertion, while dashed arrows ($--\rightarrow$) denote ply drop. (a) Valid propagation utilizing only insertions. (b) Valid propagation mixing insertion and drop with a different start sequence. (c) Invalid cyclic propagation.

Crucially, the chosen propagation path must not leave physically adjacent panels disconnected. As shown in Fig. 3,

paths (a) and (b) are valid because they systematically enforce compatibility rules across physical boundaries. In contrast, path (c) is invalid: the tree branches out, leaving adjacent panels (e.g., the 11-ply and 19-ply laminates) without a direct parent-to-child relationship. Without this explicit evaluation across their shared boundary, the generative model cannot guarantee their final compatibility [7].

2.3 From Stochastic Search to Generative Inference

Although the SPD method simplifies the global topology into a sequence of local problems, assigning the free variables typically relies on iterative stochastic solvers. To overcome this computational latency, recent works have proposed replacing heuristic searches with Deep Generative Models.

Conditional GANs (CGAN) [6] were proposed as a Physics-Informed approach to generate sequences from mechanical targets. However, as noted by Xie et al. [8], their stochastic nature lacks interpretability. Furthermore, adversarial training suffers from inherent instabilities and is prone to *mode collapse*, leading to a severe loss of design diversity when scaling to larger sequences.

Transformer-Based Architectures were concurrently investigated by Xie et al. [8], who introduced *LayupFormer* to formulate sequence retrieval as a linguistic translation task. While promising, their implementation relies on explicit grammar masking applied *after* the Softmax layer and a complex hybrid loss function. Computing this loss requires training and maintaining a high-fidelity differentiable surrogate, introducing significant architectural overhead and complex hyperparameter tuning. Crucially, these generative models operate solely at the single-laminate level, ignoring inter-laminate blending constraints.

Our proposed neuro-symbolic approach, in contrast, explicitly decouples physical performance from topological constraints. The neural agent, SeqGPT, approximates the inverse distribution $P(S \mid \xi^*)$ to inherently satisfy local constraints (MC1–MC5) without intermediate surrogates. To address the global topology, our Constrained Beam Search practically implements the SPD strategy to enforce blending (MC0). By treating inter-panel derivation as a constrained tree search, symbolic logic strictly prunes invalid ply drops and insertions. Throughout this process, SeqGPT acts as a physics-driven heuristic to evaluate and rank the valid branches, transforming an intractable combinatorial problem into a tractable, beam-limited exploration.

3 Methodology

The methodology is presented from four perspectives: (1) The definition of a domain-specific tokenization strategy that maps physical laminate properties to discrete sequence tokens, (2) The architecture of the SeqGPT agent, (3) The dataset generation and training details, and (4) The neuro-symbolic decoding protocol for multi-panel generation.

3.1 AI Modeling and Tokenization Strategy

As introduced in Section 2.1, the generative model is conditioned on the continuous target vector ξ^* and the target laminate thickness N . To integrate these physical targets into the neural architecture, ξ^* and N are concatenated and directly mapped into the model’s latent space via a linear projection layer.

The Symmetry constraint (MC1) optimizes the generation process: the neural agent predicts only the first half of the sequence ($n = \lceil N/2 \rceil$ plies). The full laminate is then reconstructed by exact mirroring, with the final generated token acting as the central axis for odd number of plies.

Next-token prediction. Standard angles $\{0^\circ, 90^\circ, \pm 45^\circ\}$ are mapped to integer tokens $\{1, 2, 3, 4\}$, alongside [PAD] (0) and [EOS] (5) tokens for training. During inference, since the target thickness N is provided in the prompt, generation stops deterministically once the required n plies are produced.

Initialization Prompt. To ensure outer surface protection (Damage Tolerance), we enforce a fixed boundary condition where all generated laminates must begin with a $[45^\circ, -45^\circ]$ block, mapped to tokens $[3, 4]$. In our Decoder-Only architecture, the autoregressive generation is triggered by an initial context acting as a prompt and denoted as [BOS]. This initial state, [BOS], encapsulates the physical conditioning vector $[PHYS] = (\xi^*, N)$ alongside these fixed boundary tokens. This initializes the sequence as [BOS] = $[PHYS, 3, 4]$, ensuring that every generated layup is physically conditioned from the first step and satisfies the mandatory surface requirements by construction.

3.2 Architecture and Input Conditioning

We employ a conditional Transformer architecture based on the **minGPT** framework [3] (see Figure 4). The model is a lightweight Transformer Decoder composed of **3 layers** with **3 attention heads**. The embedding dimension is set to $d_{model} = 48$, resulting in approximately 83k trainable parameters. Following the standard minGPT architecture, each transformer block incorporates Layer Normalization, Masked Multi-Head Attention, and a Multilayer Perceptron (MLP) featuring a GELU non-linear activation function. Dropout layers are also included to act as a regularization technique and prevent overfitting.

We have two different input types: conditional input parameters $[PHYS] = (\xi^*, N)$ and discrete ply tokens $\{1, 2, 3, 4\}$. To mix them, we project them into a shared space of dimension d_{model} .

Because Transformers require fixed-length input vectors while optimal laminates exhibit variable thicknesses, we standardize the sequence capacity based on our dataset’s maximum thickness, $N_{max} = 41$. Consequently, we define a maximum context window of $n_{max} = \lceil N_{max}/2 \rceil = 21$. During inference, SeqGPT autoregressively generates the half-sequence until n ply tokens have been produced. Any remaining sequence length up to n_{max} is subsequently filled with [PAD] tokens.

Feature Fusion and Positional Encoding. We use two projection layers to enter the common latent space:




Figure 4: Architecture of the SeqGPT agent. The model is a lightweight Decoder-Only Transformer.

- The **conditional input** vector of physical targets is processed by a Linear Layer. This projects floating-point values into vectors $v_{phys} \in \mathbb{R}^{d_{model}}$.
- The **discrete tokens** pass through a standard trainable Embedding Layer. This maps each integer to a dense vector $v_{ply} \in \mathbb{R}^{d_{model}}$.

We concatenate these vectors to form the input sequence, with the physical parameters prepended as a conditioning prompt for generation. Because transformers treat tokens as unordered, this architecture adds learned positional embeddings to the sequence vectors. This allows the model to capture the exact positional context of each token within the stack.

Attention Mechanism and Physical Alignment. During the self-attention calculation, every new token generated at step t attends back to all previous positions $\theta_{<t}$. Crucially, it also attends to the initial positions corresponding to v_{phys} . The attention weights implicitly learn to correlate specific ply orientations to match the physical target.

Training Objective and Masking. The model operates autoregressively: at each step it predicts the probability distribution of the next token conditioned on the sequence history and the physical context. The network is trained by minimizing the Cross-Entropy loss. However, the conditional input token [PHYS] and the padding token [0] must be excluded from the optimization. To achieve this, a binary mask m_t is introduced such that $m_t = 0$ for positions corresponding to physical parameters or padding tokens, and $m_t = 1$ for valid ply tokens. For a sequence with maximum length n_{max} , the masked Cross-Entropy loss over the vocabulary space V is defined as

$$\mathcal{L} = - \sum_{t=1}^{n_{max}} m_t \sum_{j=1}^{|V|} \hat{y}_{j,t} \log(y_{j,t}), \quad (2)$$

where $|V|$ denotes the vocabulary size, $\hat{y}_{j,t} \in \{0, 1\}$ is the one-hot encoded ground-truth token for class j at position t , and

$$y_{j,t} = P(\theta_t = j \mid \theta_{<t}, \xi^*, N)$$

is the probability predicted by the softmax layer for token j at step t , conditioned on the previous plies and the target physical parameters. The objective optimized during training corresponds to the average of this loss over all tokens in a mini-batch.

3.3 Dataset and Training Details

The model was evaluated after being trained on a synthetic dataset of 20,000 unique, valid stacking sequences. The dataset was generated via an Acceptance-Rejection Monte Carlo sampling strategy: an extensive pool of over 10^6 randomly generated discrete layups was evaluated, and only those strictly satisfying the local manufacturing rules (MC1-MC5) were retained. For each valid sequence, the corresponding Nemeth parameters ($\alpha, \beta, \gamma, \delta$) were analytically computed using the Classical Laminite Theory (CLT) based on IMA/M21E material properties ($E_1 = 154$ GPa, $E_2 = 8.5$ GPa, $G_{12} = 4.2$ GPa, $\nu_{12} = 0.35$).

ID	Constraint	Description / Condition
C1	Sequence Heritage	Drop: Inherit plies from S_p sequentially ($S_c \subset S_p$). The total number of skipped parent plies must not exceed $N_p - N_c$. Insertion: Inherit all plies from S_p sequentially ($S_p \subset S_c$). The total number of newly generated plies must not exceed $N_c - N_p$.
C2	Balance	S_p must provide counter-angles to offset current S_c imbalance.
C3	Proportion	Sufficient angles of each orientation remain to satisfy the ply proportion constraint.

the log-probability score. These penalties bias the beam ranking without affecting absolute feasibility. For instance, during a "drop" mode, a continuity penalty favors inheriting contiguous plies from the parent rather than creating large structural gaps.

Furthermore, to guarantee robustness, a **Conservative Fallback Strategy** is introduced. If the constrained beam search fails to produce a valid sequence for a target thickness N_{target} (often due to overly restrictive constraints), the target thickness is iteratively increased. In the worst case, the parent stacking sequence is fully copied ($N_{\text{target}} = N_{\text{parent}}$). This ensures that a manufacturable solution is always returned.

4 Numerical Results and Discussion

This section presents the evaluation of the proposed framework, moving from intrinsic model assessment to industrial validation. Section 4.1 first characterizes the generative performance independently of the beam search solver, evaluating the model's capacity to strictly respect manufacturing constraints, accurately target buckling parameters, and provide interpretability through self-attention mechanisms. Section 4.2 then validates the full neuro-symbolic framework on the standard 18-panel horseshoe benchmark. Finally, Section 4.3 discusses these results by comparison to the baseline in terms of mechanical margins and computational efficiency.

4.1 Single Panel Generation with SeqGPT

4.1.1 Diversity-Validity Trade-off under Top- k Sampling

We evaluated the generative behavior of SeqGPT independently of the deterministic beam-search solver by sampling 25 sequences for each of 200 buckling targets. Diversity was controlled through Top- k sampling, with $k = 1$ corresponding to greedy decoding and $k > 1$ introducing stochastic exploration. No generated sequence appeared in the training set, indicating that the model learns physical correlations rather than memorizing configurations.

Figure 6 shows that while the model captures most industrial rules, increasing k degrades validity. Given the small vocabulary size ($|V| = 4$), setting $k = 4$ effectively removes any pruning of the output distribution. As a result, the sampler is forced to consider all possible ply orientations, including low-probability but physically invalid transitions (e.g., $0^\circ \rightarrow 90^\circ$), which leads to occasional violations of the disorientation rule.

A moderate setting ($k = 2$) provides a trade-off between diversity and feasibility, but strict manufacturing compliance without external repair requires greedy decoding ($k = 1$).

4.1.2 Accuracy of Mechanical Target Conditioning

An effective inverse design framework must accurately recover the continuous physical properties prescribed by the macroscopic optimization phase. We evaluated the constrained model capacity to target specific parameters ($\alpha, \beta, \gamma, \delta$) on 1000 unseen target vectors using greedy decoding ($k = 1$).




Figure 6: Evolution of the sequence validity rate (left axis) and the breakdown of specific constraint violations (right axis) as a function of the top- K sampling parameter.




Figure 7: Cumulative distribution of the Mean Normalized Error (MNE) across 1000 generated sequences. The dashed line defines the 15.0% industrial tolerance threshold.

Exact matching of the continuous targets is often impossible due to the discrete nature of ply angles ($\{0^\circ, \pm 45^\circ, 90^\circ\}$) and the cubic dependence of flexural stiffness. Consequently, targeting performance is evaluated using the Mean Normalized Error (MNE). Figure 7 displays the cumulative distribution of the targeting error. Using a 15% MNE industrial tolerance threshold, the framework achieves an overall success rate of **81.0%**.

These controlled error rates confirm that the latent conditioning mechanism guides the generative process toward the required mechanical regions. SeqGPT then acts as an efficient heuristic locator that guides the beam search presented in the previous section.

4.1.3 Model Interpretability: Self-Attention Analysis

To understand the model’s decision-making process, we analyzed the multi-head self-attention weights from the final transformer layer. These weights represent the network’s focus on previously generated plies and conditioning targets during autoregressive decoding.

The visualizations in Fig. 8 suggest that the attention heads might capture different aspects of the sequence. As expected from causal masking, all matrices exhibit a lower-triangular structure. Head 2 (Fig. 8b) shows a strong concentration of attention weights along the immediate sub-diagonal. This suggests a primary focus on the directly preceding ply, which is consistent with the model learning local Markovian constraints (such as contiguity or maximum disorientation). In contrast, Head 1 (Fig. 8a) captures longer-range dependencies across the entire sequence history, which could indicate an aggregation of global context, potentially to monitor cumulative properties like the overall $\pm 45^\circ$ balance.

Physical Alignment. The first column of both matrices shows how the model attends to the initial conditioning token ([PHYS]). Head 1 maintains persistent attention on this token, ensuring continuous grounding in the physical targets. Conversely, Head 2’s attention to [PHYS] progressively decays as the sequence lengthens. Towards the end of the generation, local ply-to-ply constraints (the sub-diagonal) become so dominant that the model focuses almost exclusively on the preceding ply and ignores both the distant sequence history and the initial physical token. The third attention head is omitted due to redundancy.

4.2 Multi-Panel Generation: 18-Panel Horseshoe Benchmark

4.2.1 Test Case

The performance of the proposed neuro-symbolic approach is evaluated on the reference 18-panel horseshoe benchmark ($r = 18$) [2]. As shown in Figure 9, this problem consists of interconnected panels with fixed biaxial loadings and dimensions. The structural material is Graphite/Epoxy IM7/8552 ($E_1 = 141$ GPa, $E_2 = 9.03$ GPa, $G_{12} = 4.27$ GPa, $\nu_{12} = 0.32$, ply thickness 0.191 mm). Panel thickness is bounded between $n_{min} = 14$ and $n_{max} = 48$ plies. While the original benchmark by Irisarri et al. [2] explores an extended orientation set (steps of 15°),




Figure 8: Self-attention weight distributions from the final layer of the SeqGPT model. (a) Head 1: Global context. (b) Head 2: Local Markovian constraints.

we restrict our design space to the industrial standard set $\{0^\circ, \pm 45^\circ, 90^\circ\}$. This choice aligns with conventional manufacturing certifications and focuses the challenge on achieving the target nemeth parameters with a more constrained, discrete vocabulary. The objective is to minimize total mass while ensuring a buckling Reserve Factor $RF > 1$. We retrieve optimal thickness distributions and buckling parameters directly from [2] to use as physical targets for sequence generation.




Fig. 8. Eighteen-panel test problem [8,26], all loads in lbf/in ($\times 175.1$ for N/m).


Figure 9: Eighteen-panel test problem benchmark layout [2]. All local loadings N_x and N_y are in lbf/in.

4.2.2 Propagation Map Definition

To translate the local geometries and biaxial loadings defined in the test case (Fig. 9) into manufacturable laminates, the multi-panel optimization process follows the standard bi-level strategy introduced in Section 2.1.

The first step of this framework—the continuous optimization phase—takes these geometric boundaries and internal loads (N_x, N_y) as inputs to minimize the total structural mass. This continuous relaxation yields the optimal target vector ξ^* alongside the target thickness N_{target} for each individual panel i . This output can be formally defined as the set of continuous physical targets $\mathcal{T} = \{(\xi_i^*, N_{target,i})\}_{i=1}^{18}$.

Because the core focus of our study is the discrete stacking sequence recovery, we bypass executing the continuous FEA ourselves. Instead, we directly extract this global optimum target set $\mathcal{T}$ from the reference solution provided by Irisarri et al. [2].



4.2.3 Comparative Analysis: SeqGPT vs. Evolutionary Algorithms

$$\lambda_{(m,n)} = \frac{\pi^2 \left[D_{11} \left(\frac{m}{a} \right)^4 + 2(D_{12} + 2D_{66}) \left(\frac{m}{a} \right)^2 \left(\frac{n}{b} \right)^2 + D_{22} \left(\frac{n}{b} \right)^4 \right]}{\left(\frac{m}{a} \right)^2 N_x + \left(\frac{n}{b} \right)^2 N_y}. \quad (4)$$

posed constructive approach dynamically prunes invalid sequences during the beam search, generating the entire set of blended sequences in a single inference pass. This reduces the total recovery time to under **4.5 seconds** (with a beam width of 10).

3. **System Maintainability:** The explicit decoupling between the data-driven physical model (neural) and the manufacturing rules (symbolic) provides a major advantage for the system's lifecycle. While the symbolic decoder is primarily mandated to enforce the global blending constraint, it can readily be used to filter local rules as well. If industrial manufacturing constraints evolve, these symbolic filters can be updated directly, without requiring costly dataset regeneration or model retraining.

Conclusion

This paper introduced SeqGPT, a neuro-symbolic framework for the inverse design of blended composite structures under manufacturing constraints. By reformulating stacking sequence retrieval as a conditional sequence generation problem and combining a lightweight Transformer with constrained decoding, the approach eliminates the main computational bottleneck of classical bi-step composite optimization workflows.

The key contribution lies in the explicit decoupling between data-driven modeling of mechanical correlations and symbolic enforcement of blending rules, ensuring manufacturable solutions by construction. Results on the 18-panel horseshoe benchmark demonstrate buckling performances comparable to state-of-the-art evolutionary approaches, with recovery times reduced from hours to seconds.

Several directions for future work are identified to address current limitations and extend the framework. First, the current formulation aggregates all physical parameters into a single context embedding. Future developments will investigate more structured conditioning schemes based on multiple context tokens, enabling finer control of the self-attention mechanism and improving the alignment between continuous performance targets and discrete stacking sequence realizability. Second, multi-panel composite structures are naturally represented as graphs. While the current Search Propagation Direction (SPD) is manually defined through template-based strategies, future work will explore systematic graph-based formulations of propagation rules. Such an approach would enable the automatic inference of propagation paths for complex topologies while preserving the guarantees provided by symbolic constrained decoding.

Acknowledgments

This research was supported by Airbus Operations SAS. This work was carried out as part of the ModIA apprenticeship program (INSA Toulouse and Toulouse INP-ENSEEIH).

Generative AI Usage

In accordance with the European Commission's recommendations, the authors acknowledge the use of generative AI (Google Gemini) for English language revision and assistance with TikZ figure generation. The authors assume full responsibility for all scientific content, methodologies, and conclusions presented in this manuscript.

References

- [1] J. E. Herencia, P. M. Weaver, and M. I. Friswell. Optimization of long anisotropic laminated fiber composite panels with t-shape stiffeners. *AIAA Journal*, 45(10):2497–2510, 2007.
- [2] François-Xavier Irisarri, Agathe Lasseigne, Fabien Simon, and Rodolphe Le Riche. Evolutionary optimization of stacking sequence tables for composite structures with blending and ply-drop constraints. *Journal of Composite Materials*, 48(26):3285–3303, 2014.
- [3] Andrej Karpathy. mingpt: A pytorch reimplementation of gpt. <https://github.com/karpathy/minGPT>, 2020.
- [4] Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, J. Jung, Chandra Bhagavatula, and Yejin Choi. Neurologic decoding: (un)supervised neural text generation with lexical constraints. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4288–4299, 2021.
- [5] Michael P Nemeth. Importance of anisotropy on buckling of compression-loaded symmetric composite plates. *AIAA Journal*, 24(11):1831–1835, 1986.
- [6] Cheng Qiu, Yuzi Han, Logesh Shanmugam, Ying Zhao, Shaotong Dong, Shanyi Du, and Jinglei Yang. A deep learning-based composite design strategy for efficient selection of material and layup sequences from a given database. *Composites Science and Technology*, 230:109154, 2022.
- [7] Marco Scardaoni and Marco Montemurro. A general-purpose search propagation direction strategy for the global blending of composite laminates. *Composite Structures*, 299:116082, 2022.
- [8] Wenjie Xie, Hasan Caglar, Kailun Deng, David Ayre, and Yifan Zhao. Layupformer: A deep generative model for composite laminate layup design. *Composites Science and Technology*, 275:111490, 2026.