

De novo design of DNA origami with a generative diffusion model

Corresponding Author: Professor Do-Nyun Kim

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Attachments originally included by the reviewers as part of their assessment can be found at the end of this file.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

Please refer to the attached pdf.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

(Remarks on code availability)

Reviewer #3

(Remarks to the Author)

In this manuscript, Chien Truong-Quoc and colleagues present Generative SNUPI, an integrated DNA nanostructure design framework combining a diffusion-based generative model, SNUPI structural simulation, and automated sequence design, enabling de novo construction of complex DNA origami. To address the scarcity of training data, the authors generate a dataset of 450 equilibrium structures using SNUPI. By incorporating conditional diffusion sampling guided by Wasserstein Distance (WD)-based optimal transport conditioning, the model can translate line-based target geometries into physically plausible 3D base-pair coordinates. Together with a routing and sequence-design module, the authors establish an end-to-end pipeline for creating sophisticated free-form DNA architectures. Experimentally, the authors validate not only static monomers but also dynamically reconfigurable structures and modular higher-order assemblies. As someone who routinely designs DNA nanostructures using traditional CAD tools, I find this new computational approach exciting and potentially transformative.

Below are several suggestions from the perspective of DNA nanostructure design that I believe will help further strengthen the manuscript.

1. The size of the training dataset is crucial for diffusion models. The manuscript uses 450 structures with 500 noising steps each. It remains unclear whether this dataset is sufficiently large and diverse to ensure stable training, robust learning of structural priors, and reliable generalization. Additional justification or analysis (e.g., convergence behavior, validation performance, or ablation studies) would strengthen confidence in the model's reliability.
2. From the assembly results, Generative SNUPI appears to struggle with structures containing sharp turns, such as Mona Lisa, the Pagoda, and the second design in Fig. 3b. These regions often exhibit distortion or reduced fidelity. It would be valuable to discuss whether the authors have explored more effective strategies to construct or stabilize such sharp-turn geometries. Furthermore, the general applicability of the method to structures with high curvature or abrupt angular features

requires additional evidence or clarification, e.g. Design ID 10.

3. For the "hexa-chiral" and "rotating triangle" designs, the observed deformation appears to arise from insufficient rigidity at the junctions. Could local stiffening—such as incorporating 6HB edges or other multi-helix reinforcements—mitigate this deformation and preserve the intended mechanical behavior?

(Remarks on code availability)

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

I appreciate the care to which the authors addressed our concerns.
The resubmitted paper is now suitable for publication as is.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

(Remarks on code availability)

I reviewed the code and found it acceptable.

Reviewer #3

(Remarks to the Author)

This manuscript presents an interesting generative framework for DNA origami design by combining diffusion-based geometry generation with deterministic routing and sequence assignment. I strongly agree with Reviewer 1's concern regarding the actual contribution of the model relative to traditional algorithms. In particular, the revised manuscript makes it clearer that the AI component primarily generates base-pair-level geometries, whereas topology generation, routing, and sequence assignment still rely on downstream deterministic procedures.

The additional analyses addressing memorization and mode collapse are helpful and partially reduce concerns regarding the limited training set. The authors also now describe the current scope and limitations more clearly, including the dependence on rule-based post-processing and the restricted applicability to the present structural class.

While the evidence for broad generalization remains somewhat limited, I find the work technically valuable and suitable for publication as long as the claims remain appropriately restrained.

(Remarks on code availability)

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Detailed Responses to the Reviewers' Comments

Manuscript ID: NCOMMS-25-88809-T

Manuscript Type: Article

Title: “**De novo design of DNA origami with a generative diffusion model**”

Author(s): Chien Truong-Quoc, Kyoungghwa Jeon, Jinho Kim, Seo Hyun Kwon, Dongsik Seo, Chanseok Lee, Do-Nyun Kim

We appreciate the reviewers' thorough reading of our manuscript and useful comments. Our point-by-point response follows.

Reviewer #1 (Remarks to the Author):

1. Summary

The authors present “Generative SNUPI,” a novel generative diffusion-based design framework, which can generate complex DNA origami nanostructures in free-form. By shifting from rule-based to data-driven design, this work expands the accessible design space to include arbitrary wireframe architectures, biomimetic shapes, and reconfigurable mechanical systems. I have significant concerns regarding the computational novelty, the validity of the dataset, and the actual contribution of the AI model versus traditional algorithms. The current evidence suggests the diffusion model acts primarily as a coordinate point-cloud generator, while the complex biochemical topology relies on deterministic post-processing and manual intervention. Once the statements around the major findings of the paper are adjusted to this more limited, but yet still sufficiently novel, claim, and my concerns are addressed, the paper should be appropriate for publication.

2. Major Concerns

Q1. The authors claim to have used a generative diffusion model, but their training dataset only contains 450 individual wireframe structures (216 2D and 234 3D designs). For a diffusion model intended to learn complex structural distributions, this dataset is extremely small. Of concern is the possibility that the model could simply be learning the noise and not the signal. Crucially, the authors concur with this concern in SI Note S5 where they note that the unconditional model struggles with large structures and often fails due to the small and structurally simple dataset. Supplementary Fig. S10 suggests the model has not learned the underlying physical rules of DNA origami but is likely overfitting to the training shapes. The authors should provide a rigorous analysis of overfitting. For example, they could calculate and present the geometric similarity (e.g.,

RMSD) between the generated structures and their nearest neighbors in the training set to prove the model is not simply retrieving memorized examples.

► We thank the reviewer for raising this important concern. We first clarify why a feasible generative model can be learned despite the limited number of training structures (~450). Next, we performed two additional evaluations to test whether the model (i) retrieve training examples or (ii) collapse to deterministic outputs under strong conditioning.

Our model is trained using a standard diffusion objective to predict injected noise across multiple noise levels under a fixed forward noising process. Then, sampling proceeds via iterative denoising from Gaussian noise. This approach fundamentally differs from deterministic reconstruction, as the model is not trained to directly reproduce specific input structures.

In our setting, several factors further make learning feasible despite the finite size of training set. First, the task is conditional: generation is guided by an input target geometry, which substantially narrows the space of plausible designs compared to unconstrained distribution in arbitrary 3D space. Second, a common DNA origami structure contains a large number of base pairs up to ~7,000 (Supplementary Note S3), providing a rich collection of node level local geometric contexts for message passing and learning even when the number of structures is limited. Third, the SE(3)-equivariant graph architecture imposes a strong inductive bias (roto-translation equivariance), reducing sensitivity to absolute coordinates and supporting generalization to target shapes.

While these considerations support why effective learning is feasible despite the limited size of the training set, we additionally performed explicit evaluations. Specifically, we tested whether the model retrieves training examples or collapses to deterministic outputs under strong conditioning. These analyses confirm that the model generates genuinely novel structures and maintains stochastic sampling behavior rather than memorizing or deterministically reproducing training data (Supplementary Fig. S4–S5).

Experiment 1: Unconditional sampling (anti-memorization test)

We generated $n = 100$ unconditional samples with 400–500 base-pair range (Supplementary Fig. S4a). For each generated sample G_i , we selected 30 training structures whose base-pair counts are closest to $N(G_i)$. We then performed rigid alignment using the iterative closest point (ICP) algorithm and computed mean RMSD (in nm) against each candidate. The nearest-neighbor distance for G_i is defined as the minimum RMSD across the 30 comparisons:

$$d_i = \min_{T \in C_{30}(G_i)} \text{RMSD}_{ICP}(G_i, T)$$

Results. Supplementary Fig. S4b shows eight representative generated structures spanning diverse geometries from planar (2D-like) to out-of-plane (3D) configurations, together with their nearest neighbors from the generated set. The nearest-neighbor RMSD distribution does not

collapse near zero, with a mean of 1.07 nm and a median of 1.02 nm across the 100 generated samples (Supplementary Fig. S4c). If the generated structures were simple copies of training structures, we would expect a non-trivial fraction of samples to exhibit near-zero nearest-neighbor RMSD.

We further report the RMSD distributions for representative target families (triangle, square, pentagon) within the range of 400-500 base pairs, which remain clearly separated from zero (Supplementary Fig. S4a, S4d). The means are 1.81 nm, 2.14 nm, and 2.63 nm, respectively—values that are even larger than the mean of the nearest-neighbor RMSD within the generated structures. Together, these results support that the model generates significant structural variability rather than retrieving training examples.

Experiment 2: Fixed target shape with fixed number of base pairs experiment (anti-collapse / non-determinism test)


To test whether the conditional model collapses under strong constraints, we fixed both the target shape and the number of base pairs and generated $n = 100$ designs with different random seeds. The RMSD distribution between individual generated structures and the target shape remains clearly separated from zero (Supplementary Fig. S5), and overlay visualization shows stochastic dispersion around the target rather than a single identical configuration, indicating that the model is not deterministic and does not collapse under fixed conditioning.


► Supplementary Fig. S4 and Supplementary Fig. S5 were added to supplementary materials. We also clarified the paragraph in the section of “Generative SNUPI design framework” in the Manuscript as follows.

- From: These results suggest that the framework successfully accommodates common DNA origami structures across a wide range of sizes, achieving high fidelity in shape generation even for structures consisting of approximately 15,000 base pairs (Supplementary Fig. S3). (...)
- To: These results suggest that the framework successfully accommodates common DNA origami structures across a wide range of sizes, achieving high fidelity in shape generation even for structures consisting of approximately 15,000 base pairs (Supplementary Fig. S3).

In addition, we assessed novelty and sampling variability using two complementary analyses to evaluate potential overfitting. For unconditional sampling of 400-500 base-pair structures, nearest-neighbor root mean square deviation (RMSD) evaluations within the generated samples and RMSD comparisons to the training set showed that generated structures are sufficiently distinct from the training data (< 1.0% of samples below 0.5 nm RMSD; Supplementary Fig. S4). Furthermore, the conditional sampling with both target shape and the number of base pairs fixed confirmed non-deterministic sampling without collapse despite strong conditioning (Supplementary Fig. S5). These results demonstrate that Generative SNUPI learns a generative prior rather than simply reproducing samples from the training set and remains non-deterministic under target

shape conditioning without mode collapse. (...)





Q2. The manuscript claims “De novo design,” implying the AI generates the DNA sequences and topology. However, the diffusion model only outputs 3D coordinates, while the critical DNA topology—including scaffold routing, spanning tree construction, and staple assignment—is handled by a deterministic routing algorithm. Since the biochemical validity is ensured by the post-processing algorithm, not the diffusion generative model, the claim that the diffusion model “designs” the DNA origami appears to be overstated. The AI appears to function as a 3D meshing tool. The authors should clarify the distinction between the AI’s output (geometry) and the routing algorithm’s output (topology).

► We thank the reviewer for the thoughtful comment and agree that it is important to clearly distinguish the roles of the diffusion model and the deterministic routing algorithm. We also appreciate the opportunity to clarify the meaning of “de novo design” in the context of this work and have revised the manuscript accordingly to avoid any possible overstatement.

In our framework, the diffusion model generates a base-pair-level 3D structural realization, whereas the routing module deterministically produces a biochemically valid topology and sequence set (e.g., loop conversion, spanning tree construction, scaffold routing, and staple assignment). The routing step is therefore essential for finalizing biochemical validity.

Although the diffusion model outputs coordinates, these are not generic geometric vertices. Both the training data and generated samples are represented as base-pair coordinate point clouds, where each point corresponds to a physical base-pair step. In particular, the coordinates encode biophysically meaningful constraints, including (i) DNA helical features reflecting the local geometry, and (ii) an approximately constant inter-base-pair rise consistent with dsDNA. Thus, the diffusion output provides physically grounded geometric constraints at the base-pair scale.

Therefore, biochemical validity is not independent of the diffusion output: several routing steps directly use the diffusion-generated base-pair geometry. For example, the locations of loop double-crossovers are determined by selecting spatially adjacent base pairs near edge centers, and the orientations of remaining base pairs are then propagated from these crossover sites (Supplementary Note S5, Steps 3 and 6). In this sense, while routing is deterministic, it depends on and is constrained by the feasibility of the diffusion-generated base-pair-level geometry.

In this work, “de novo design” refers to the generation of a new base-pair-level structural realization at each sampling, even under the same target shape condition. Because routing is determined by the generated geometry rather than a pre-implemented canonical template, different diffusion samples lead to distinct crossover patterns and topological realizations.

To make this point explicit, we added Supplementary Fig. S6, which analyzes topology diversity induced by repeated sampling:


- Panels a–c summarize topology statistics for 100 generated triangle designs (the 100 structures shown in Supplementary Fig. S5; both target shape and base-pair count fixed). The overlap circular map and crossover-count histogram show non-collapsing crossover distributions across samples under identical conditioning.
- Panels d–f show three routed designs for Mona Lisa_v1 (~7000-base scaffold), illustrating clearly different crossover patterns despite using the same target shape.

These additions clarify that our framework implements a one-to-many mapping from a fixed target shape to multiple feasible DNA origami designs: the diffusion model provides stochastic and physically constrained base-pair geometries, and the routing module converts each geometry into a DNA valid topology and sequence design.

► We added Supplementary Fig. S6 and clarified the paragraph in the section of “Generative

SNUPI design framework” in the Manuscript as follows.

- From: (...) In the final sequence design step, nucleotide sequences are assigned to both scaffold and staple strands, with staple lengths ranging from 20 to 60 nucleotides, ensuring compatibility with experimental synthesis constraints. (...)
- To: (...) In the final sequence design step, nucleotide sequences are assigned to both scaffold and staple strands, with staple lengths ranging from 20 to 60 nucleotides, ensuring compatibility with experimental synthesis constraints. **Importantly, while the routing procedure is deterministic given the generated geometry, the diffusion model produces a different base-pair-level structural geometry at each sampling. Consequently, repeated generations under the same target shape can induce distinct routed crossover and topology patterns (Supplementary Fig. S6). (...)**



Supplementary Figure S6. Topology diversity from repeated generation under fixed target shape.

a, Representative circular map of staple crossover topology for one representative design from generated structure in Supplementary Fig. S5. **b**, Overlaid circular map of 100 generated designs for the triangle target shape shown in Supplementary Fig. S5 (both target shape and the number of base pairs fixed), showing the diversity of crossover distributions. **c**, Histogram of the number of staple crossovers per design across the 100 designs. **d–f**, Three representative routed designs for Mona Lisa_v1 using a standard ~7000-base scaffold,

illustrating distinct crossover patterns despite the same target shape. Link color in the circular maps indicates crossover distance along the scaffold: blue, short-range; red, long-range (coolwarm colormap).

3. Minor Concerns

Q1. The manuscript highlights the design of auxetic and transformable structures. However, according to the SI, the flexible hinges required for rotation were created by selectively cutting unpaired staple segments manually. The text should explicitly disclose the extent of manual intervention required to achieve the functional states (hinges and connectors) and avoid attributing these engineered features solely to the generative capability of the model.

► In response to the reviewer's comments, we have clarified the extent of manual intervention required to achieve the shape transition. While the generative model defines the overall geometry of the transformable structures, the formation of flexible hinges requires a deliberate post-design modification. Specifically, we now describe how engineered defects are introduced at the junction by removing unpaired bases to create single-stranded gaps. This manual adjustment is intended to reduce local stiffness and increase junction flexibility, thereby promoting monomer open-to-close shape transition and suppressing unintended multimer formation.

► We clarified the paragraph in the section of "Design of transformable structures and modular assemblies using Generative SNUPI" in the Manuscript as follows.

- From: (...) To promote open-to-close shape transition, we introduced engineered defects at the junctions via single-stranded gaps⁵³ proven to increase the junction flexibility and prevent the formation of multimers (Supplementary Note S8). (...)
- To: (...) To promote open-to-closed shape transitions, [we introduced short single-stranded gaps⁶³ at the junctions after shape generation \(Supplementary Note S9\). These manually modified gaps reduce local stiffness, thereby increasing the junction flexibility and preventing the formation of multimers.](#) (...)

Q2. While Generative SNUPI establishes a robust foundation for generative modeling, it is currently restricted to wireframe architectures composed of 2HB edges. Extending the framework to encode more rigid components, such as six-helix bundles (6HB) or multilayer lattices, requires overcoming the current scarcity of standardized training datasets.

► As the reviewer points out, extending Generative SNUPI beyond 2HB wireframe nanostructures is currently limited by the availability of standardized training datasets. We have revised the manuscript to clarify both the challenge and potential strategies to address this limitation. For 2D and 3D 6HB-edge structures, we note that training datasets can be constructed in a manner similar to this study by combining established wireframe design tool ATHENA (including TALOS and METIS) with SNUPI simulations to extract base-pair coordinates for target shape. In contrast, multilayer lattice-based structures remain more challenging due to the small number of reported designs currently available. To address this limitation, we suggest that incorporating data augmentation strategies will be necessary to expand the effective training dataset and enable the generation of a broader range of rigid structural motifs.

► We clarified the paragraph in the section of “Discussion” in the Manuscript as follows.

- From: (...) However, the flexibility of the framework allows further expansion to 6HB-edge systems, multilayer lattices, and sheet-like geometries as sufficient training datasets become available. Integration with deep learning–based analysis tools such as DeepSNUPI³⁴ is expected to play a key role in accelerating the design and validation of new structures by enabling rapid and accurate predictions of structural and functional behaviors. (...)
- To: (...) However, the flexibility of the framework allows further extension to 6HB-edge systems, multilayer lattices, and sheet-like geometries as sufficient training datasets become available. While a substantial amount of 6HB wireframe structural data can be obtained from related studies^{25,58} combined with SNUPI simulations, multilayer lattice designs remain scarce due to the limited number of reports. Incorporating data augmentation strategies could help address this limitation by expanding the training dataset and enabling the generation of a broader range of structural variations. Integration with deep learning–based analysis tools such as DeepSNUPI⁴¹ is expected to further accelerate the design and validation of new structures by enabling rapid and accurate predictions of structural and functional behaviors.

Q3. Please clarify the exact definition of “pseudo-atomic” in this context. Does the final output file (e.g., PDB) contain all-atom coordinates generated via idealized helical parameters, or does it remain a coarse-grained representation (e.g., backbone beads only)? If the model is all-atom, please briefly describe the geometric algorithm used to reconstruct the full atomic details (orientation, major/minor grooves) from the coarse-grained diffusion output.

► We thank the reviewer for the clarification request. To avoid potential confusion, we have replaced the term “pseudo-atomic” with “atomic representation” throughout the revised manuscript. Here, “atomic” refers to an exported PDB file in which full atomic coordinates are generated by a deterministic, template-based reconstruction during post-processing (CNDO → oxDNA → PDB), rather than predicted by the diffusion model. We provide the conversion details below for clarity.

In our framework, the diffusion model outputs a base-pair-level coarse geometry, and local orientations are assigned later during deterministic routing. After routing, we export the final design into a CanDo.cndo representation, which stores strand topology (scaffold/staples) and base-pair frames via dNode(i) (base-pair center) and triad(i) (base-pair reference frame).

We then use tacoxDNA to convert CNDO into oxDNA, and further convert oxDNA nucleotide rigid-body frames into an atomic PDB using a template-based geometric reconstruction (i.e., the atoms are placed according to standard base/sugar/phosphate templates consistent with the underlying coarse-grained nucleotide frames).

► We also clarified the paragraph in the “Generative SNUPI design framework” section of the

Manuscript and added a reference to tacoxDNA as follows.





- From: (...) In the final sequence design step, nucleotide sequences are assigned to both scaffold and staple strands, with staple lengths ranging from 20 to 60 nucleotides, ensuring compatibility with experimental synthesis constraints. This process enables the efficient and precise translation of generated DNA nanostructures into experimentally realizable constructs, facilitating the conversion of 100 generated designs into pseudo-atomic models with staple sequences for validation.
- To: (...) In the final sequence design step, nucleotide sequences are assigned to both scaffold and staple strands, with staple lengths ranging from 20 to 60 nucleotides, ensuring compatibility with experimental synthesis constraints. Importantly, while the routing procedure is deterministic given the generated geometry, the diffusion model produces a different base-pair-level structural geometry at each sampling. Consequently, repeated generations under the same target shape can induce distinct routed crossover and topology patterns (Supplementary Fig. S6). **We export the final design into a CanDo⁵⁹ representation and then convert it into an atomic PDB representation via a deterministic post-processing pipeline using tacoxDNA⁶⁰ (CNDO → oxDNA → PDB). Atomic models generated through this process enable further validation through experimental realization and computational analysis.**

Q4. The manuscript highlights the limitations of conventional algorithmic approaches. However, a direct quantitative comparison with state-of-the-art automated design tools (e.g., DAEDALUS (Veneziano et al.,[5] Science 2016)) is missing. Please provide a comparative analysis, ideally in a tabular format benchmarking Generative SNUPI against these established methods.

► Following the reviewer’s suggestion, we have updated a direct comparative analysis between Generative SNUPI and the established automated wireframe design tools (DAEDALUS and PERDIX), summarized in a tabular format in the revised supplementary materials (Supplementary Table S1). This comparison benchmarks key design capabilities such as allowable edge length and vertex angle variability, crossover placement, and support for curved edges. As shown in the comparison, while conventional algorithmic approaches rely on predefined geometric constraints, Generative SNUPI enables greater design freedom by allowing arbitrary features including curved edges through its diffusion-based generative process. This comparison clarifies the complementary strengths of Generative SNUPI relative to existing methods.

► Supplementary Table S1 is added.

	DAEDALUS [25]	PERDIX [14]	This study
Design method	Deterministic (Spanning tree algorithm)	Deterministic (Rendering dual graph and spanning tree algorithm)	Combined generative diffusion model and deterministic routing algorithm(similar to PERDIX)

Edge type	2HB	2HB	2HB
Input geometry	3D wireframe polyhedron	2D wireframe	2D or 3D wireframe
Edge constraints	Multiples of 10.5 bps	Arbitrary edge length	Arbitrary edge length
Vertex angles	Uniform	Variable	Variable
Crossover regularity	Regular	Regular	Irregular
Edge curvature allowance	No	No	Yes
Example (2D) (Similar scaffold length)	N/A	Circle-like design (Segmented edge)  Scaffold length = 744nt	Circular design (Curved edge)  Scaffold length = 758nt
Example (3D) (Similar scaffold length)	Sphere-like design (Segmented edge)  Scaffold length = 5279nt	N/A	Sphere design (Curved edge)  Scaffold length = 4919nt
Input file type	PLY	PLY	SVG, PLY, STL, PT
Output file type	Sequence list (CSV), CanDo, BILD, PDB, caDNAno(JSON)	Sequence list (CSV), CanDo, BILD, PDB, caDNAno(JSON)	CanDo, BILD, oxDNA, PDB, PyTorch tensors
User interfaces	Command line based	Command line based	Command line based

4. Suggested Citations

There are several key references that provide context and prior art to this work that should be cited:

1. While the manuscript cites Jun et al. (ATHENA) and Zhang et al., it appears to miss the seminal work that established the “top-down” routing of arbitrary 3D wireframe meshes using tension propagation.[5]

2. The manuscript highlights SNUPI as a rapid simulation alternative to Molecular Dynamics (MD). While oxDNA is covered via Doye et al., there are other "rapid" simulation tools that serve as relevant benchmarks or related work. [4]

3. The manuscript utilizes SE(3)-Transformers for 3D roto-translation equivariance. The methodology section would benefit from acknowledging the broader Geometric Deep Learning field, specifically efficient alternatives that influenced the current wave of diffusion models. [2, 9]

4. To better position this work within the rapidly evolving research landscape, I recommend that the authors briefly acknowledge concurrent advances in ML-driven DNA nanotechnology: recent studies on generative design [3, 6], predictive modeling [10, 1], and computational characterization[7, 8].

References

[1] Jared Huzar, Roxana Coreas, Markita P. Landry, and Grigory Tikhomirov. Aibased prediction of protein corona composition on dna nanostructures. *ACS Nano*, 19(4):4333–4345, 2025. PMID: 39772513.

[2] Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24240–24253. Curran Associates, Inc., 2022.

[3] Zhi Lin. Algebraic and generative design of dna origami via origami monoids. *bioRxiv*, 2025.

[4] Victor Garcia Satorras, Emiel Hoogetboom, and Max Welling. E(n) equivariant graph neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9323–9332. PMLR, 18–24 Jul 2021.

[5] Remi Veneziano, Sakul Ratanalert, Kaiming Zhang, Fei Zhang, Hao Yan, Wah Chiu, and Mark Bathe. Designer nanoscale dna assemblies programmed from the top down. *Science*, 352(6293):1534–1534, 2016.

[6] Anthony J Vetturini, Jonathan Cagan, and Rebecca E Taylor. Generative designenabled exploration of wireframe dna origami nanostructures. *Nucleic Acids Research*, 53(2):gkae1268, 12 2024.

[7] Yuchen Wang, Xin Jin, and Carlos Castro. Accelerating the characterization of dynamic dna origami devices with deep neural networks. *Scientific Reports*, 13(1):15196, sep 2023.

[8] Xingfei Wei, Qiankun Mo, Chi Chen, Mark Bathe, and Rigoberto Hernandez. Characterizing dna origami nanostructures in tem images using convolutional neural networks. *Journal of*

Chemical Information and Modeling, 65(13):6526–6536, 2025. PMID: 40540709.

[9] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In International Conference on Learning Representations, 2022.

[10] Judith Zubia-Aranburu, Andrea Gardin, Lars Paffen, Matteo Tollemeto, Ane Alberdi, Maite Termenon, Francesca Grisoni, and Tania Patino Padial. Predicting dna origami stability in physiological media by machine learning. bioRxiv, 2025.

► We appreciate the reviewer for pointing out relevant prior studies that were not sufficiently cited in the original manuscript. We have added the suggested references [28-29, 36-40, 54] to the corresponding sections to provide a more comprehensive overview including CAD and simulation tools, diffusion model for designing small molecules, and machine learning-driven DNA nanotechnology.

► We enriched the paragraph in the section of “Introduction” in the Manuscript as follows.

- From: Despite the proliferation of computer-aided design (CAD) tools^{21–27}, most frameworks remain constrained by algorithmic regularity and periodic crossover placement, following the inherent geometry of double-stranded DNA. (...) Recently, diffusion-based generative AI models³¹ have demonstrated astonishing performance in the inverse design and generation of complex protein structures^{32,33}. By learning implicit structural priors from data, these models offer a data-driven alternative to conventional algorithmic approaches. However, developing a data-driven AI model for designing DNA nanostructures is challenging due to the lack of a sufficiently large and standardized dataset of their designs suitable for training generative models³⁴ and highly limited structures determined either experimentally or computationally^{35–40}.
- To: Despite the proliferation of computer-aided design (CAD) tools^{21–29}, most frameworks remain constrained by algorithmic regularity and periodic crossover placement, following the inherent geometry of double-stranded DNA. (...) Recently, diffusion-based generative AI models³³ have demonstrated astonishing performance in the inverse design and generation of [structures including complex protein structures^{34,35} and small molecules^{36,37}](#). By learning implicit structural priors from data, these models offer a data-driven alternative to conventional algorithmic approaches. [However, despite recent advances in machine learning-based predictive modeling³⁸ and computational characterization of DNA nanostructures^{39,40}, developing a generative AI model for designing DNA nanostructures remains challenging. This difficulty primarily arises from the lack of a sufficiently large and standardized dataset of their designs suitable for training generative models⁴¹ and highly limited structures determined either experimentally or computationally^{42–47}.](#)

► We included the additional reference [54] in the section of “Generative SNUPI design framework” in the Manuscript as follows.

- From: (...) We incorporate SE(3) symmetry-based geometric invariance^{45,46} to directly predict base-pair coordinates from noisy inputs, minimizing the mean-squared error (MSE) between the predicted and ground-truth structures without requiring any alignment³² (Fig. 2b and Supplementary Note S2). This loss function ensures denoising paths match the data distribution, leading to constructible DNA origami structures.
- To: (...) We incorporate SE(3) symmetry-based geometric invariance⁵²⁻⁵⁴ to directly predict base-pair coordinates from noisy inputs, minimizing the mean-squared error (MSE) between the predicted and ground-truth structures without requiring any alignment³⁴ (Fig. 2b and Supplementary Note S2). This loss function ensures denoising paths match the data distribution, leading to constructible DNA origami structures.

Reviewer #3 (Remarks to the Author):

In this manuscript, Chien Truong-Quoc and colleagues present Generative SNUPI, an integrated DNA nanostructure design framework combining a diffusion-based generative model, SNUPI structural simulation, and automated sequence design, enabling de novo construction of complex DNA origami. To address the scarcity of training data, the authors generate a dataset of 450 equilibrium structures using SNUPI. By incorporating conditional diffusion sampling guided by Wasserstein Distance (WD)-based optimal transport conditioning, the model can translate line-based target geometries into physically plausible 3D base-pair coordinates. Together with a routing and sequence-design module, the authors establish an end-to-end pipeline for creating sophisticated free-form DNA architectures. Experimentally, the authors validate not only static monomers but also dynamically reconfigurable structures and modular higher-order assemblies. As someone who routinely designs DNA nanostructures using traditional CAD tools, I find this new computational approach exciting and potentially transformative.

Below are several suggestions from the perspective of DNA nanostructure design that I believe will help further strengthen the manuscript.

Q1. The size of the training dataset is crucial for diffusion models. The manuscript uses 450 structures with 500 noising steps each. It remains unclear whether this dataset is sufficiently large and diverse to ensure stable training, robust learning of structural priors, and reliable generalization. Additional justification or analysis (e.g., convergence behavior, validation performance, or ablation studies) would strengthen confidence in the model's reliability.

► We thank the reviewer for raising this important concern. We first clarify why a feasible generative model can be learned despite the limited number of training structures (~450). Next, we performed two additional evaluations to test whether the model (i) mimic training examples or (ii) collapse to deterministic outputs under strong conditioning.

Our model is trained using a standard diffusion objective to predict injected noise across multiple noise levels under a fixed forward noising process. Then, sampling proceeds via iterative denoising from Gaussian noise. This approach fundamentally differs from deterministic reconstruction, as the model is not trained to directly reproduce specific input structures.

In our setting, several factors further make model reliable despite the finite size of training set. First, the task is conditional: generation is guided by an input target geometry, which substantially narrows the space of plausible designs compared to unconstrained distribution in arbitrary 3D space. Second, a common DNA origami structure contains a large number of base pairs up to ~7,000 (Supplementary Note S3), providing a rich collection of node level local geometric contexts for message passing and learning even when the number of structures is limited. Third, the SE(3)-equivariant graph architecture imposes a strong inductive bias (roto-translation equivariance), reducing sensitivity to absolute coordinates and supporting generalization to target shapes.

While these considerations support why effective learning is feasible despite the limited size of

the training set, we additionally performed explicit evaluations. Specifically, we tested whether the model reproduces training examples or collapses to deterministic outputs under strong conditioning. These analyses confirm that the model generates genuinely novel structures and maintains stochastic sampling behavior rather than memorizing or deterministically reproducing training data (Supplementary Fig. S4–S5).

Experiment 1: Unconditional sampling (anti-memorization test)

We generated $n = 100$ unconditional samples with 400–500 base-pair range (Supplementary Fig. S4a). For each generated sample G_i , we selected 30 training structures whose base-pair counts are closest to $N(G_i)$. We then performed rigid alignment using the iterative closest point (ICP) algorithm and computed mean RMSD (in nm) against each candidate. The nearest-neighbor distance for G_i is defined as the minimum RMSD across the 30 comparisons:

$$d_i = \min_{T \in C_{30}(G_i)} \text{RMSD}_{\text{ICP}}(G_i, T)$$

Results. Supplementary Fig. S4b shows eight representative generated structures spanning diverse geometries from planar (2D-like) to out-of-plane (3D) configurations, together with their nearest neighbors from the generated set. The nearest-neighbor RMSD distribution does not converge near zero, with a mean of 1.07 nm and a median of 1.02 nm across the 100 generated samples (Supplementary Fig. S4c). If the generated structures were simple copies of training structures, we would expect a non-trivial fraction of samples to exhibit near-zero nearest-neighbor RMSD.

We further report the RMSD distributions for representative target families (triangle, square, pentagon) within the range of 400–500 base-pairs, which remain clearly separated from zero (Supplementary Fig. S4a, S4d). The means are 1.81 nm, 2.14 nm, and 2.63 nm, respectively—values that are even larger than the mean of the nearest-neighbor RMSD within the generated structures. Together, these results support that the model generates significant structural variability rather than memorizing training examples.

Experiment 2: Fixed target shape with fixed number of base pairs experiment (anti-collapse / non-determinism test)


To test whether the conditional model collapses under strong constraints, we fixed both the target shape and the number of base pairs and generated $n = 100$ designs with different random seeds. The RMSD distribution between individual generated structures and the target shape remains clearly separated from zero (Supplementary Fig. S5), and overlay visualization shows stochastic dispersion around the target rather than a single identical configuration, indicating that the model is not deterministic and does not collapse under fixed conditioning.


► Supplementary Fig. S4 and Supplementary Fig. S5 were added to supplementary materials. We also clarified the paragraph in the section of “Generative SNUPI design framework” in the

Manuscript as follows.

- From: These results suggest that the framework successfully accommodates common DNA origami structures across a wide range of sizes, achieving high fidelity in shape generation even for structures consisting of approximately 15,000 base pairs (Supplementary Fig. S3). (...)
- To: These results suggest that the framework successfully accommodates common DNA origami structures across a wide range of sizes, achieving high fidelity in shape generation even for structures consisting of approximately 15,000 base pairs (Supplementary Fig. S3).

In addition, we assessed novelty and sampling variability using two complementary analyses to evaluate potential overfitting. For unconditional sampling of 400-500 base-pair structures, nearest-neighbor root mean square deviation (RMSD) evaluations within the generated samples and RMSD comparisons to the training set showed that generated structures are sufficiently distinct from the training data (< 1.0% of samples below 0.5 nm RMSD; Supplementary Fig. S4). Furthermore, the conditional sampling with both target shape and the number of base pairs fixed confirmed non-deterministic sampling without collapse despite strong conditioning (Supplementary Fig. S5). These results demonstrate that Generative SNUPI learns a generative prior rather than simply reproducing samples from the training set and remains non-deterministic under target shape conditioning without mode collapse. (...)






Q2. From the assembly results, Generative SNUPI appears to struggle with structures containing sharp turns, such as Mona Lisa, the Pagoda, and the second design in Fig. 3b. These regions often exhibit distortion or reduced fidelity. It would be valuable to discuss whether the authors have explored more effective strategies to construct or stabilize such sharp-turn geometries. Furthermore, the general applicability of the method to structures with high curvature or abrupt angular features requires additional evidence or clarification, e.g. Design ID 10.

► We thank the reviewer for this important point. In the current version of our framework, we

indeed observed that geometries with abrupt angular features are more prone to distortion, as observed in AFM imaging. This limitation is primarily attributed to (i) electrostatic repulsion between closely positioned DNA edges and (ii) flexible junction compliance, which can destabilize sharp-turn regions unless additional geometric/mechanical constraints are incorporated. Regarding structures with high curvature, we note that the outer rings of hexa-chiral design (Fig. 4b vi) exhibit high fidelity to the target shape. This suggests that high curvature alone does not necessarily limit shape integrity. Instead, our observations indicate that pronounced local fluctuations play a more critical role in destabilizing curved structures such as Design ID 10. These local fluctuations may arise from the large discrepancy between the Euclidean distance between two vertices and the prescribed edge length.

To provide quantitative evidence and practical guidance, we performed a controlled set of RMSD analysis to examine how the length or radius of curved edges and interior angles affect the deviation between the designed configuration and the predicted equilibrium configuration. As summarized in Supplementary Note S7, the RMSD distributions (min–mean–max) show a systematic increasing trend as the angular geometry becomes sharper and the curved-edge segment becomes longer.



► While a comprehensive exploration of stabilization strategies (including automated feasibility warnings and extensive parametric/experimental validation) is an important direction for future

work, it is beyond the scope of the present manuscript. Therefore, we do not add a full section to the main text, but instead (i) include the results as Supplementary Note S7 and (ii) add a brief clarification in the section of “Experimental validation of Generative SNUPI designs” in the Manuscript as follows:

- From: (...) To reduce structural deviation, we enhanced the stiffness the design by adding edges to the flexible regions, resulting in the Face 2 design (Fig. 4a, bottom row). AFM images for the Face 2 design demonstrated significantly reduced structural distortion with a slightly increased monomer yield of 83.1% (Fig. 4a, bottom row; Supplementary Fig. S5). The RMSD distribution quantitatively confirmed an improved match with the target shape (3.45 ± 0.35 nm). While manually done here, this revision process could be automated by integrating Generative SNUPI with a typical optimization algorithm owing to its predictive capability.
- To: (...) To reduce structural deviation, we enhanced the stiffness the design by adding edges to the flexible regions, resulting in the Face 2 design (Fig. 4a, bottom row). AFM images for the Face 2 design demonstrated significantly reduced structural distortion with a slightly increased monomer yield of 83.1% (Fig. 4a, bottom row; Supplementary Fig. S8). The RMSD distribution quantitatively confirmed an improved match with the target shape (3.45 ± 0.35 nm). *In addition to local flexibility, extreme geometric features, such as vertices with sharp turns, can further increase structural deviations from the target shape due to electrostatic repulsion (Supplementary Note S7). Although performed manually in this study, this revision process could be automated by integrating Generative SNUPI with an optimization algorithm, given its predictive capability.*

Q3. For the "hexa-chiral" and "rotating triangle" designs, the observed deformation appears to arise from insufficient rigidity at the junctions. Could local stiffening—such as incorporating 6HB edges or other multi-helix reinforcements—mitigate this deformation and preserve the intended mechanical behavior?

► We thank the reviewer for the insightful suggestion regarding local stiffening strategies to mitigate junction deformation. Increasing local rigidity at the junctions can reduce deformation in intermediate states. However, such stiffening often compromises the intended transformable behavior of the structures, which relies on sufficient junction flexibility.

The primary goal of the kirigami-based auxetic metastructures was to demonstrate the versatility of our framework in generating diverse transformable structures rather than to optimize mechanical rigidity. In practice, we manually modified the junctions to increase their flexibility in order to achieve the fully closed state. Prior to these manual interventions, the structures preferentially formed multimers rather than undergoing a monomeric shape transition (Supplementary Figs. S17-S18). This observation highlights an inherent trade-off between junction reinforcement and reversible shape transformation, and in the present study we prioritized reconfigurability to demonstrate the geometric fidelity offered by Generative SNUPI.

Comments to the Authors

1. Summary

The authors present “Generative SNUPI,” a novel generative diffusion-based design framework, which can generate complex DNA origami nanostructures in free-form. By shifting from rule-based to data-driven design, this work expands the accessible design space to include arbitrary wireframe architectures, biomimetic shapes, and reconfigurable mechanical systems. I have significant concerns regarding the computational novelty, the validity of the dataset, and the actual contribution of the AI model versus traditional algorithms. The current evidence suggests the diffusion model acts primarily as a coordinate point-cloud generator, while the complex biochemical topology relies on deterministic post-processing and manual intervention. Once the statements around the major findings of the paper are adjusted to this more limited, but yet still sufficiently novel, claim, and my concerns are addressed, the paper should be appropriate for publication.

2. Major Concerns

1. The authors claim to have used a generative diffusion model, but their training dataset only contains 450 individual wireframe structures (216 2D and 234 3D designs). For a diffusion model intended to learn complex structural distributions, this dataset is extremely small. Of concern is the possibility that the model could simply be learning the noise and not the signal. Crucially, the authors concur with this concern in SI Note S5 where they note that the unconditional model struggles with large structures and often fails due to the small and structurally simple dataset. Supplementary Fig. S10 suggests the model has not learned the underlying physical rules of DNA origami but is likely overfitting to the training shapes. The authors should provide a rigorous analysis of overfitting. For example, they could calculate and present the geometric similarity (e.g., RMSD) between the generated structures and their nearest neighbors in the training set to prove the model is not simply retrieving memorized examples.
2. The manuscript claims “De novo design,” implying the AI generates the DNA sequences and topology. However, the diffusion model only outputs 3D coordinates, while the critical DNA topology—including scaffold routing, spanning tree construction, and staple assignment—is handled by a deterministic routing algorithm. Since the biochemical validity is ensured by the post-processing algorithm, not the diffusion generative model, the claim that the diffusion model “designs” the DNA origami appears to be overstated. The AI appears to function as a 3D meshing tool. The authors should clarify the distinction between the AI’s output (geometry) and the routing algorithm’s output (topology).

3. Minor Concerns

1. The manuscript highlights the design of auxetic and transformable structures. However, according to the SI, the flexible hinges required for rotation were created by selectively cutting unpaired staple segments manually. The text should explicitly disclose the extent of manual intervention required to achieve the functional states (hinges and connectors) and avoid attributing these engineered features solely to the generative capability of the model.
2. While Generative SNUPI establishes a robust foundation for generative modeling, it is currently restricted to wireframe architectures composed of 2HB edges. Extending the framework to encode more rigid components, such as six-helix bundles (6HB) or multilayer lattices, requires overcoming the current scarcity of standardized training datasets.
3. Please clarify the exact definition of "pseudo-atomic" in this context. Does the final output file (e.g., PDB) contain all-atom coordinates generated via idealized helical parameters, or does it remain a coarse-grained representation (e.g., backbone beads only)? If the model is all-atom, please briefly describe the geometric algorithm used to reconstruct the full atomic details (orientation, major/minor grooves) from the coarse-grained diffusion output.
4. The manuscript highlights the limitations of conventional algorithmic approaches. However, a direct quantitative comparison with state-of-the-art automated design tools (e.g., DAEDALUS (Veneziano et al., [5] Science 2016)) is missing. Please provide a comparative analysis, ideally in a tabular format benchmarking Generative SNUPI against these established methods.

4. Suggested Citations

There are several key references that provide context and prior art to this work that should be cited:

1. While the manuscript cites Jun et al. (ATHENA) and Zhang et al., it appears to miss the seminal work that established the "top-down" routing of arbitrary 3D wireframe meshes using tension propagation.[5]
2. The manuscript highlights SNUPI as a rapid simulation alternative to Molecular Dynamics (MD). While oxDNA is covered via Doye et al., there are other "rapid" simulation tools that serve as relevant benchmarks or related work. [4]
3. The manuscript utilizes SE(3)-Transformers for 3D roto-translation equivariance. The methodology section would benefit from acknowledging the broader Geometric Deep Learning field, specifically efficient alternatives that influenced the current wave of diffusion models. [2, 9]
4. To better position this work within the rapidly evolving research landscape, I recommend that the authors briefly acknowledge concurrent advances in ML-driven DNA nanotechnology: recent studies on generative design [3, 6], predictive modeling [10, 1], and computational characterization[7, 8].

References

- [1] Jared Huzar, Roxana Coreas, Markita P. Landry, and Grigory Tikhomirov. Ai-based prediction of protein corona composition on dna nanostructures. *ACS Nano*, 19(4):4333–4345, 2025. PMID: 39772513.
- [2] Bowen Jing, Gabriele Corso, Jeffrey Chang, Regina Barzilay, and Tommi Jaakkola. Torsional diffusion for molecular conformer generation. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24240–24253. Curran Associates, Inc., 2022.
- [3] Zhi Lin. Algebraic and generative design of dna origami via origami monoids. *bioRxiv*, 2025.
- [4] Víctor Garcia Satorras, Emiel Hoogetboom, and Max Welling. E(n) equivariant graph neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 9323–9332. PMLR, 18–24 Jul 2021.
- [5] Rémi Veneziano, Sakul Ratanaalert, Kaiming Zhang, Fei Zhang, Hao Yan, Wah Chiu, and Mark Bathe. Designer nanoscale dna assemblies programmed from the top down. *Science*, 352(6293):1534–1534, 2016.
- [6] Anthony J Vetturini, Jonathan Cagan, and Rebecca E Taylor. Generative design-enabled exploration of wireframe dna origami nanostructures. *Nucleic Acids Research*, 53(2):gkae1268, 12 2024.
- [7] Yuchen Wang, Xin Jin, and Carlos Castro. Accelerating the characterization of dynamic dna origami devices with deep neural networks. *Scientific Reports*, 13(1):15196, sep 2023.
- [8] Xingfei Wei, Qiankun Mo, Chi Chen, Mark Bathe, and Rigoberto Hernandez. Characterizing dna origami nanostructures in tem images using convolutional neural networks. *Journal of Chemical Information and Modeling*, 65(13):6526–6536, 2025. PMID: 40540709.
- [9] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. In *International Conference on Learning Representations*, 2022.
- [10] Judith Zubia-Aranburu, Andrea Gardin, Lars Paffen, Matteo Tollemeto, Ane Alberdi, Maite Termenon, Francesca Grisoni, and Tania Patiño Padial. Predicting dna origami stability in physiological media by machine learning. *bioRxiv*, 2025.