CONSTRUCTAL LAW CONFERENCE – CLC2024
Design in Nature and Evolution
Romanian Academy, Bucharest, 10–11 October 2024

Proceedings of CLC2024
pp. 41–44

CONSTRUCTAL LAW, BIOMIMICRY, AND TOPOLOGY OPTIMIZATION THROUGH THE LENS OF GENERATIVE AI

MATEI C. IGNUTA-CIUNCANU, PHILIP TABOR, RICARDO F. MARTINEZ-BOTAS

SETTL Lab, Imperial College, London, SW7 2AZ
mi3618@ic.ac.uk, philip.tabor2@imperial.ac.uk, r.botas@imperial.ac.uk

Efficient heat management in the context of semiconductor miniaturization demands unconventional solutions to navigate constrained design spaces. The emergence of unsupervised learning as a versatile design companion capable of processing multi-disciplinary datasets presents promising avenues for streamlined bio-inspired solutions for fundamental conduction problems. As a case study, we tackle the pivotal area-to-point problem that led to the formulation of the Constructal Law.

This study proposes a novel methodology harnessing *denoising diffusion probabilistic models* (DDPM) to seamlessly integrate features from topology optimization, constructal theory, and biomimetic structures. We use DDPM as a digital non-equilibrium complex system to generate geometric patterns for conductive heat sink design. These patterns are sampled to address the original formulation of the area-to-point problem.

This study emphasizes the significance of morphological freedom in addressing multi-objective problems, reinforcing arguments first formulated within the constructal paradigm. Our framework facilitates low-cost sampling of intricate shapes capable of serving multiple temperature objectives by synthesizing principles from different fields. We show the emergence of dendritic structures to solve distribution problems in an unsupervised learning scenario, drawing a parallel between information and energy flows.

This research underscores the transformative potential of Generative AI in blending design features across disparate disciplines, a potent tool for developing conductive heat sink solutions beyond deterministic optimization approaches.

Keywords: Generative AI; Morphological freedom; Thermal management; Nature-inspired design.

1. INTRODUCTION

The exponential increase in computing power as part of the digital transformation follows a direct consequence of a two-order of magnitude reduction in transistor scale. Formulated in 1965, Moore's Law accurately predicted doubling the number of transistors in an integrated circuit every two years. The more stringent requirements for heat dissipation capacity and the ever-smaller scale size progressively ruled out natural convection, fan, and liquid cooling. They ultimately limited the space available for high-conductivity material [1].

The cooling of electronic devices has become a hierarchical problem, considering the different objectives at each scale. At the chip level, the goal is to maintain a minimal temperature despite the high local fluxes. At the module level, where the heat loading becomes prominent, a minimal temperature gradient becomes relevant to reduce losses. Nonetheless, most consumer electronics operate under extensive variations in ambient conditions, making the ability to handle dynamic loads a priority at the system level.

Initially studied for electronics cooling applications, the area-to-point problem [2] contributed to formulating the Constructal Law as a general principle for resource allocation. This problem addresses bottlenecks in confined spaces and has seen a variety of alternative solutions, from gradient-based topology optimization through simulated bionic growth to particle swarm optimization [3–5].

The applicability of the Constructal Law as a design tool for challenges across various design spaces demonstrates that solutions developed for the ATP problem can be scaled seamlessly. Additionally, technologies like additive manufacturing open the design space, requiring new computational solutions to explore non-conventional avenues of inquiry.

2. COMPUTATIONAL METHODOLOGY

In this context, we propose a unifying generative approach to draw parallels between living and non-living systems to explore effective bi-material patterns. For this purpose, we use a diffusion algorithm developed to replicate non-equilibrium thermodynamics [6] for generating de novo structures. This method allows design feature blending and fast unsupervised sampling from a multi-modal heterogeneous dataset. We show that the

DOI: https://doi.org/10.59277/CLC/2024.09

AI agent can navigate a combinatorial design space and that the workflow enhances the morphological freedom of the samples, creating de novo candidates that visually display a high level of hybridity.

In the context of generative AI, diffusion models produce a Markov chain that gradually maps a known distribution (*e.g.*, a Gaussian) onto the underlying distribution of the dataset. Replicating statistical mechanics processes, the algorithm learns to estimate the steps for the inverse probability flow required to reach the smooth target distribution [7]. This strategy was motivated by the authors considering the Kolmogorov forward and backward equations, which showed that the two paths can be covered by respective small-perturbation diffusion processes that take the same functional form [6].

In the annealed importance sampling configuration, the algorithm estimates the multi-variate probability distribution from a single sample taken from the forward trajectory. Simulating a quasi-static physical process, the generative model assumes infinitesimally low diffusion rates and estimates the probability distributions of the intermediary steps using the Jarzynski equality [8].

In thermodynamics, this equality relates the exponential average of work done during non-equilibrium transformations to the free energy difference between two states. The denoising probabilistic model fundamentally resembles Langevin dynamics in discretizing the stochastic Fokker-Planck equation [7]. Hence, the amount of energy (here, information) lost along the path between the equilibrium distributions is proportional to the step size. This parallel provides a qualitative understanding of the diffusion rate as the critical parameter controlling sample quality [9].

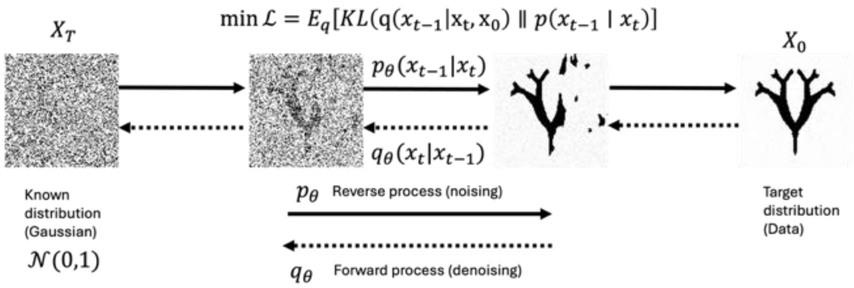

Fig. 1 – Diffusion probabilistic model that uses Markov chains that add (reverse process) and remove (forward process) noise and minimize local cross-entropy losses to establish computational and information efficient paths between the data and Gaussian equilibrium distributions.

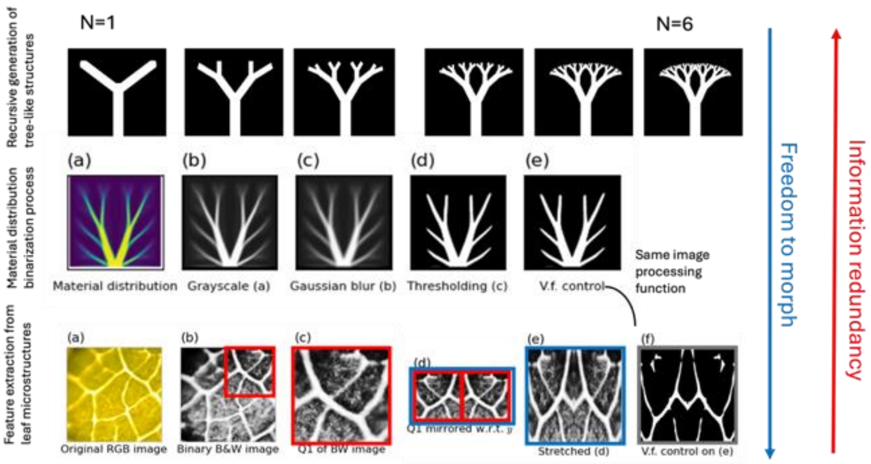

Fig. 2 – Three-modal dataset: hierarchical tree-like patterns with nested splitting, pre-optimized samples with boundary conditions and objectives to match the ATP problem, and bio-inspired leaf motifs with self-similarity. The freedom to morph and the degree of *orderliness* of the three families of samples are inversely correlated.

Given the method's unsupervised learning nature, the input data set will dominate the sampled structures. For the two-dimensional area-to-point (ATP) problem, we have curated a collection of microstructures: digital dendritic patterns with recursive scaling, output samples from an in-house topology optimizer following solid isotropic material with penalization, and processed close-up leaf pictures available online [10].

We reproduced the steps to normalize the three data modalities for brevity to compile a cohesive data set in Fig. 2. 11k bi-material samples represent each category. The recursive trees, which have a prescribed shape and an *orderly* appearance, carry the least amount of information. Conversely, the leaf structures contain self-similar and local symmetry-breaking features and require information-dense descriptions.

3. **RESULTS**

Figure 3 reproduces the results of unconditioned diffusion from the data set. They show that the solid inductive biases embedded in the three modalities can generate plausible candidates without an objective function feedback loop. Contrasting supervised learning aimed at ATP temperature objectives, the diffusion model's high-resolution synthesis improved feature interbreeding and produced some asymmetric samples through the open-ended generative process. The samples' diversity and the improved freedom to morph can be used for designing bi-material solutions for multi-objective problems with arbitrary boundary conditions.

From a constructal perspective, the non-equilibrium diffusion system improves the information flow access and enlarges the communication channels between the three categories in the input data set. Nonetheless, the form-function correlation is shown to be preserved in all the samples proposed having dendritic structures. The incremental changes caused by the noising/denoising processes, controlled by (information) entropy generation minimization, are thus well calibrated for configuration and evolution toward design freedom and, ultimately, better performance. Most importantly, the diffusion framework put forward provides the information flow architecture required to bridge disparate knowledge domains, growing novel branching forms that could quickly adapt for different physics and/or objectives rather than following prescriptions, targeting an *end design*, or mimicking nature [11–13].

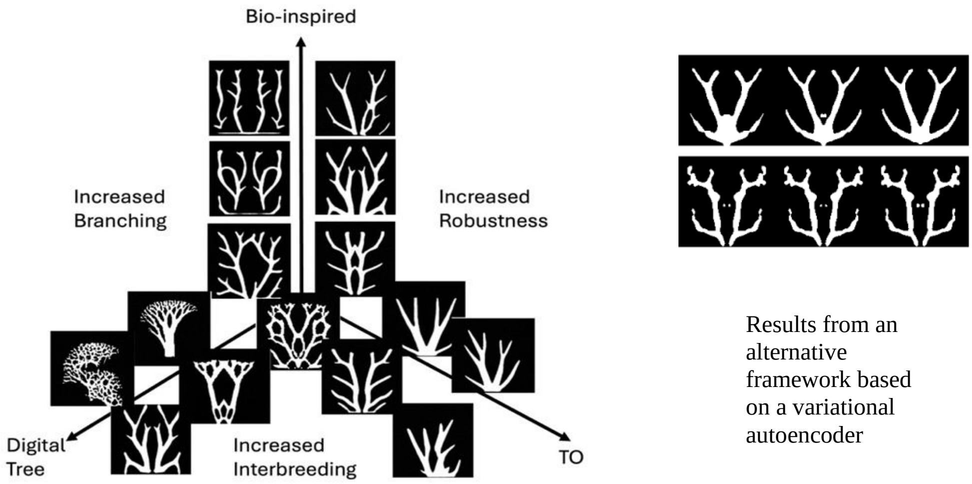

Fig. 3 – Unconditioned generation of *de novo* samples, ordered by origin and feature characteristics. The high resolution of the diffusion network yields improved mixing between design features from the three sources, showing enhanced freedom to morph compared to the samples obtained in a supervised learning setup (reproduced on the right [14]).

Moving forward from the original area-to-point problem, we have extended the capabilities of generative models beyond conductive heat sinks. The bio-inspired non-uniform structures are used to handle complex loading conditions, where temperature or heat flux varies across the surface. This adaptability is key in more advanced applications such as battery packs or high-performance computing systems where uniform cooling patterns would not be optimal. In this regard balancing thermal performance, material use, and manufacturing constraints with the hierarchical challenges posed at the beginning requires flexibility beyond conventional parametric design.

In Fig. 4, we show an example of a battery module assembly designed by the Formula Student Imperial Racing Green team. The temperature profile of the battery pack displays a clear temperature gradient. By morphing bio-inspired form to engineering function, the thermal solution adapts to the requirement of the system for enhanced performance and efficiency.

4. **CONCLUSIONS**

To conclude, we have presented a unifying generative design framework that improves the modelling resolution of thermal solutions beyond parametric design by enhancing the morphological freedom. This is made possible by the application of entropy generation minimization on the information flows to navigate complexity and produce dendritic structures from the border of chaos and order. This extends the application of constructal design to digital flow systems, demonstrating how natural flow architectures can be replicated probabilistically to optimize information and energy transfer, resulting in efficient, adaptive solutions for complex thermal and distribution challenges.

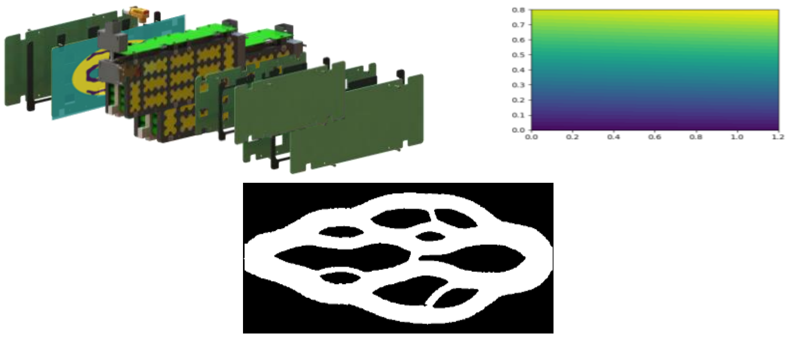

Fig. 4 – Solid model and temperature distribution of a generic battery pack (top).
Bio-inspired cooling channel topology generated from diffusion sample matching (bottom).

ACKNOWLEDGMENTS

MIC would like to thank Hannes Stark for the technical conversations on generative AI methods.

REFERENCES

1. Bejan A., Evolution in thermodynamics, *Appl. Phys. Rev.*, **4**, *1* (2017).
2. Bejan A., Constructal-theory network of conducting paths for cooling a heat generating volume, *Int. J. Heat Mass Transf.*, **40**, *4*, pp. 799–816 (1995).
3. Marck G., Nemer M., Harion J.L., Russeil S., Bougeard D., Topology optimization using the SIMP method for multiobjective conductive problems, *Numerical Heat Transfer, Part B: Fundamentals*, **61**, *6*, pp. 439–470 (2012).
4. Li B., Hong J., Yan S., Liu H., Ge L., Generating optimal heat conduction paths based on bionic growth simulation, *International Communications in Heat and Mass Transfer*, **83**, pp. 55–63 (2017).
5. Cai H., Guo K., Liu H., Xiang W., Liu C., Multiobjective optimization of area-to-point heat conduction structure using binary quantum-behaved PSO and Tchebycheff decomposition method, *Canadian Journal of Chemical Engineering*, **99**, *5*, pp. 1211–1227 (2021).
6. Sohl-Dickstein J., Weiss E.A., Maheswaranathan N., Ganguli S., Edu S., Deep Unsupervised Learning using Nonequilibrium Thermodynamics, *Forty-Second International Conference on Machine Learning, Vancouver Convention Center* (ICML), 13th – July 19th, 2025.
7. Ho J., Jain A., Abbeel P., Denoising Diffusion Probabilistic Models, *Adv. Neural Inf. Process. Syst.*, 2020, Jun.–Dec.
8. Jarzynski C., Equilibrium free energy differences from nonequilibrium measurements: a master equation approach, *Phys. Rev. E. Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics*, **56**, *5*, pp. 5018–5035 (1997).
9. Song J., Meng C., Ermon S., Denoising Diffusion Implicit Models, Cornell University, arXiv: 2010.02502.
10. Shen S.C., Buehler M.J., Nature-inspired architected materials using unsupervised deep learning, *Communications Engineering*, **1**, *1*, pp. 1–15 (2022).
11. Bejan A., Vascular flow design and predicting evolution, *International Communications in Heat and Mass Transfer*, **155**, p. 107517 (2024).
12. Bejan A., Constructal design evolution versus topology optimization, *International Communications in Heat and Mass Transfer*, **141** (2023).
13. Bejan A., Evolutionary design: Heat and fluid flow together, *International Communications in Heat and Mass Transfer*, **132**, p. 105924 (2022).
14. Ignuta-Ciuncanu M.C., Stärk H., Martinez-Botas R., *Bio-Inspired Design of Conductive Heat Sinks Using a Generative Autoencoder Framework*, https://ssrn.com/abstract=4576761.