

Folded Bodies: Topological Control of an Untrained Spiking Instrument

Lucas Kuzma

The Strange Agency

lucas@strange.agency

Abstract

We present a generative audiovisual instrument in which the performer sculpts a spiking network by folding the body it lives on. A seed point cloud is iteratively reflected across mirror planes, accumulating bilateral symmetry. The resulting cloud is both the visible body and the substrate of a leaky integrate-and-fire network: each point is a neuron, each k -nearest-neighbor edge a synapse. Geodesic farthest-point probes read local activity and drive a polyphonic sonifier with two voice paths, oscillator and granular, supporting both pitch-organized and texture-organized sonic responses to the same network state. A per-point pitch scalar derived from azimuthal position is propagated through reflections, so points related by geometric position share both color and pitch — bilateral symmetry of the body extends naturally into bilateral symmetry of the sound. Because state is addressed by position rather than index, firing patterns are spatially legible in the same view the performer uses to fold the geometry. We frame the system as a co-creative instrument whose computational-creativity contribution lies in making the morphology of the dynamical substrate itself the performer’s interface.

Introduction

Reservoir computing (RC) (Nakajima and Fischer 2021) and echo state networks (ESNs) (Jaeger and Haas 2004) offer rich dynamical substrates for music generation, but their opacity is a persistent design challenge. Kiefer (2014) demonstrates ESNs as expressive mapping layers for musical instruments while identifying *controllability* as the central difficulty: the reservoir’s state space is opaque and corresponds to no externally legible structure. Sculpting the reservoir means navigating a high-dimensional parameter space without visual or intuitive feedback, and the standard response, a trained readout layer, replaces one opacity with another. Following Kiefer, the design problem is not to eliminate the resulting *uncontrol* (Jorda 2005) — the nonlinearity and unpredictability that make a dynamical instrument expressive — but to place it where the performer can negotiate with it. Ianigro and Bown (2016) examine a similar problem for continuous-time recurrent neural networks (CTRNNs) as creative substrates and remark about the difficulty of discovery in the vastness of the CTRNN parameter space.

The system presented here proposes a different response: make the network’s topology a direct consequence of a body the performer can see and fold. The Ω operator (Bader and Oxman 2016) iteratively applies planar mirror reflections to a seed point cloud, accumulating bilateral symmetry. The resulting cloud forms the substrate of an untrained spiking network whose edges are its k -nearest-neighbor (kNN) graph. Every neuron occupies a position in a visible 3D body; its firing state at any moment is a spatial event on that body, readable at a glance. Folding the body reshapes the network, and what the performer sees folding is what they hear changing.

The system inherits its body-generating operator from a generative-morphology lineage. The L-systems descended from Lindenmayer (Prusinkiewicz and Lindenmayer 1990) show how compact rule-sets can produce bodies whose form encodes the artist’s intent; Bader and Oxman’s (2016) recursive-symmetry method introduces the planar-mirror operator Ω , applied iteratively to a seed cloud to generate complex bilaterally-symmetric forms for additive manufacturing. We adopt Ω directly, adapting it from a fabrication pipeline that produces meshes for 3D printing into a substrate-generation pipeline that produces point clouds for spiking-network construction. Hauser *et al.* (2011) formalize a related observation in soft robotics: a compliant body’s intrinsic dynamics can serve as the reservoir a recurrent system would otherwise compute on an opaque internal state. We import this framing into a software audiovisual setting where the body is a folded virtual cloud whose dynamics drive a probe-based readout— though, unlike in the standard RC pattern, our readout is not trained. The CC tradition of embodiment (Guckelsberger *et al.* 2021; Saunders *et al.* 2010) argues that situating a generative system in a visible space changes both production and reception; we do not take on its stronger production-side commitments — our body has no environment and no self-preserving dynamics — but the reception side applies directly, because bodies produced by Ω are biomorphic (bilaterally symmetric, organism-like) in a way random or force-directed renderings of the same network are not. The instrument is therefore audiovisual rather than auditory-with-visualization: the visible morphology is part of what the music is heard as coming from.




Figure 1: A folded body in performance. The same k -nearest-neighbor graph (rainbow edges) is simultaneously visible geometry, neural substrate, and probe placement domain. Geodesic probes appear as translucent spheres on the body; their activation traces scroll in the heads-up display (HUD) strip (upper left), with thin connector lines tying each probe to its trace.

Computational Creativity stance. We position *Folded Bodies* as a contribution to musical metacreation (Pasquier *et al.* 2016): the artifact is not a composition but a music-maker the performer shapes. Where standard real-time metacreation is symbolic — live coding edits text, modular synthesis re-patches topologies, evolutionary systems breed by fitness — we propose a *morphological* mode in which the maker is a body and the authorial gesture is folding. The performer chooses the body; the body chooses the music. Three concrete moves support this framing. (i) *State legibility*: because the substrate’s neurons sit in a visible body, the network’s state is a spatial event the performer reads at a glance, directly addressing the opacity problem reservoir-style instruments face (Kiefer 2014; Ianigro and Bown 2016). (ii) *Action legibility*: the same operator Ω generates the body and modifies it during performance, with no separation of design phase from performance phase. In Boden’s (2004) terms each fold acts transformationally, rewriting the space of bodies the instrument can become; we acknowledge Wiggins’s (2006) formal result that “Boden’s transformational creativity is actually exploratory creativity at the meta-level” (p. 458) and do not claim to escape that regress. (iii) *Structural coherence*: every downstream stage (kNN, probe placement, color, pitch, grain neighborhood) is geometry-indexed, so the body’s bilateral symmetry propagates through the pipeline by con-

struction — what changes visually is the same thing that changes audibly. In Kantosalo and Takala’s (2020) terms the *collective* is the performer together with the body-as-instrument, and their *contributions* divide cleanly: the performer supplies morphological intent, the system supplies dynamical interpretation, and the visible body is the shared workspace through which they negotiate.

The contribution is therefore a *system demonstration and design argument* for a morphological mode of real-time musical metacreation: the Ω operator, exposed as a live performance gesture over a geometry-indexed spiking substrate, gives the performer a legible, sculptable handle on an otherwise opaque dynamical instrument while preserving end-to-end coherence between what is seen and what is heard.

System

The instrument consists of four components: a folding operator that produces a body, a spiking network that lives on the body’s kNN graph, a set of geodesic probes that read local activity, and a polyphonic sonifier that voices each probe through selectable oscillator or granular paths. Figure 2 shows the overall flow.




Figure 2: The four components of the instrument: a folding operator that produces a body, a spiking network that lives on the body’s kNN graph, a set of geodesic probes that read local activity, and a polyphonic sonifier that voices each probe through selectable oscillator or granular paths. Arrows show information flow; colors mark the geometry (blue), dynamics (green), and sound (orange) phases.

The Ω folding operator

Starting from a seed point cloud (typically a Fibonacci-spaced sphere of N points, optionally distorted by a low-order spherical-harmonic, plane-wave noise, or Gaussian-bump field), the Ω operator applies a sequence of planar reflections. At each step, a plane is placed through the cloud. The signed half-space containing the larger subset of points is retained; the complementary half is discarded; the retained half is reflected across the plane and merged with its mirror image. The combined cloud is voxel-deduplicated and capped to N . After k steps the cloud is bilaterally symmetric with respect to the k -th plane, with earlier folds visible as bilateral substructures within each half.

A given parameter vector produces the same network topology, and small changes to plane angles produce continuous deformations of the body; see Appendix A.1 for the parameter encoding.

Spiking network on the kNN topology

A kNN graph is constructed on the final cloud. Each undirected edge becomes two directed synapses with independent forward and backward weights. The network is a leaky integrate-and-fire (LIF) system with continuous-time decay: each neuron holds a bounded membrane potential, a binary firing flag, and a refractory counter. Per tick, firing neurons distribute their weighted activation to successors; non-refractory neurons integrate inputs; neurons crossing their threshold fire and enter a refractory period during which the activation leaks quickly; multiplicative leaks are applied.

Weights, thresholds, and refractory periods are sampled from stochastic distributions keyed on a network seed, so the network has a stable “personality” that re-instantiates on any topology. The topology is a function of the folding genome; the dynamical parameters are a function of the network seed. Folding therefore reshapes the network’s connectivity graph while preserving the character of its noise — useful, because a performer can fold a body into new shapes and recognize the same instrument speaking through them. We use a spiking rather than a continuous-rate network because spike-and-refractory dynamics produce discrete, visually-legible firing events on the body and bursty, episodic activation traces that translate cleanly to musical onsets. Parameter values and the initialization kickstart are in Appendix A.2.

Geodesic probe readout

Probes (typically 5–8) are placed on the kNN graph by farthest-point sampling (FPS) using Dijkstra distance in the

graph’s intrinsic metric. Edge weights for the Dijkstra are Euclidean distances between adjacent points, so probes spread along the body’s surface even when the body is folded tightly in Euclidean space.

Each probe integrates network activation within a Euclidean neighborhood with linear distance falloff, producing a scalar activation trace per probe. The FPS algorithm and probe-slot stability across graph rebuilds are detailed in Appendix A.3.

Sonification

Each probe drives one voice. The sonifier supports two voice paths, selectable live: an *oscillator* path and a *granular* path. Probe activation drives voice gain; camera-space distance to each probe modulates filter cutoff, so closer probes sound brighter; each probe’s horizontal screen position drives stereo pan, so the audio tracks where the marker sits on screen.

Two pitch logics are exposed. In *slot mode*, pitch follows slot index, drawn from a fixed palette. In *prune mode*, pitch follows each probe’s *position on the body* via a per-point scalar that Ω carries unchanged through every fold; in the granular path the same scalar indexes a position in the source buffer rather than an oscillator frequency. Folding therefore *removes* the pitches and buffer positions the body can no longer reach, so each shape sings through a smaller, distinct palette rather than a re-shuffling of the same set; bilateral body symmetry extends to bilateral palette symmetry. See Appendix A.5.

The granular voice path was added because square oscillators, however tunable, produce a tonal output that does not match the textural quality of folded geometry. The visual surface has detail — twin clusters, density variations, fold seams — and the granular path lets the audio reflect this in kind, fusing the texture of input and the texture of output. When a clearly bilaterally-symmetric body sings in prune mode, the audio’s left and right halves are not just panned reflections but textural reflections, drawing from the same buffer regions in mirror-image patterns. This was the design goal that motivated the granular path: making what the system *sounds* like reflect what it *looks* like, not merely what it does. Voice signal chains are in Appendix A.4.

The instrument

The performer sees the body rotating in 3D with probes as colored spheres pulsing in time with their local activation, connector lines linking each probe to its on-screen HUD trace (Figure 1). Folding is a performative gesture — adding or rotating a plane reshapes both what the audience sees and

what they hear, in a way the performer can anticipate because the folded body *is* the network. Sessions typically begin by selecting a seed and plane genome, letting the network settle, adjusting neighbor count to tune activity density, and then tuning the network’s dynamics. The body constrains the search space, but can be instantly refolded with the dynamics held intact.

Demonstration

Five contrasting configurations illustrate the instrument’s range and accompany this paper as supplementary video and audio.¹

Ambient. Compact geometry with few probes and low activity, punctuated by occasional spike trains.

Expansive. Sonified folding and spike propagation across expansive geometry, with live tuning of the weight bias.

Folding. A folding sequence whose auditory behavior diversifies as the geometry grows more complex.

Spatial. Depth-driven filter modulation across a moving, expansive body, with small localized activity slowly cascading to the whole network.

Toroidal. A toroidal seed yields a sparse, knotted body; spike trains propagate slowly, producing a correspondingly sparse audio texture.

These configurations are presented as illustrative regimes rather than a formal evaluation. They are intended to demonstrate that the controls expose meaningfully distinct musical characters, and that those characters are reliably anticipated by the performer from the visible morphology.

Discussion

The system sits in a small but established ICCC neighborhood. Ianigro and Bown (2016) is the closest direct precedent: they treat continuous-time recurrent networks as creative substrates and search the parameter space for musically-affording configurations. Our move is to replace that search with a sculptable interface, exchanging exhaustive exploration for direct manipulation of a visible body whose folding is the parameter. Earlier ICCC and ICCC-adjacent work on evolutionary sonic ecosystems (McCormack 2001; Bown 2009) and embodied creative agents (Saunders *et al.* 2010) frames the broader tradition we draw on; Guckelsberger *et al.* (2021) survey the embodiment question for CC; we engage their argument from the reception vantage. The connection to morphological computation in robotics (Hauser *et al.* 2011) is more than rhetorical: a folded body whose connectivity computes is structurally the same idea expressed in a different medium.

Against the lineage of generative morphology proper — the L-systems descended from Lindenmayer (Prusinkiewicz and Lindenmayer 1990) and Bader and Oxman’s recursive-symmetry method (Bader and Oxman 2016) from which we adopt the Ω operator — the contribution is to make the resulting body a dynamical object rather than a purely sculptural one. Where their pipelines produce a shape the viewer reads or fabricates, ours produces a shape that behaves, and

whose behavior is sonifiable because both body and sonifier are indexed by the same geometry. The granular voice path is what lets this land texturally as well as structurally: a body with bilateral symmetry and visible fine detail produces sound with bilateral symmetry and matching textural detail.

Among recent ICCC demonstrations, the closest sibling is AutomaTone (Schaap and Hedblom 2024), which drives an interactive music generator from a 2D cellular automaton. The shared intuition is that a spatial process can be heard as well as seen; we differ by working in 3D morphology with a folded substrate and a spiking dynamical layer rather than a CA update rule.

Why symmetry. Because every downstream stage of the pipeline is geometry-indexed — the kNN graph, probe placement, per-point color, per-point pitch scalar, and grain neighborhood are all functions of position — the bilateral symmetry produced by Ω is preserved end-to-end by construction. Twin points (related by reflection across the final plane) therefore share color, share pitch in prune mode, and scrub the same source-buffer neighborhood in granular mode. The body’s audible behavior carries the same symmetry as the visible body as a property of the pipeline rather than as an emergent perceptual claim. Whether listeners exploit this structural symmetry as a perceptual scaffold is an open question we have not measured.

Future work. The compact, deterministic genome is built to support neuroevolutionary search: the parameter vector has the structure that indirect encodings such as Compositional Pattern Producing Networks (CPPNs) and HyperNEAT (Stanley 2007) exploit, and an evolutionary fitness landscape over musically- or visually-defined objectives is a natural extension of the sculptable-body interface presented here. We are also exploring geodesic rather than Euclidean radii for probe readout, so that fold seams act as natural separations between probe domains.

Author Contributions

The author is solely responsible for all aspects of this work.

¹<https://the.strange.agency/projects/folded-bodies>

References

- Bader, C., and Oxman, N. 2016. Recursive symmetries for geometrically complex and materially heterogeneous additive manufacturing. *Computer-Aided Design* 81:39–47.
- Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. Routledge, 2nd edition.
- Bown, O. 2009. Ecosystem models for real-time generative music: A methodology and framework. In *Proceedings of the International Computer Music Conference*.
- Clough, J., and Douthett, J. 1991. Maximally even sets. *Journal of Music Theory* 35(1):93–173.
- Guckelsberger, C.; Kantosalo, A.; Negrete-Yankelevich, S.; and Takala, T. 2021. Embodiment and computational creativity. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC)*.
- Hauser, H.; Ijspeert, A. J.; Fuchsli, R. M.; Pfeifer, R.; and Maass, W. 2011. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics* 105(5–6):355–370.
- Ianigro, S., and Bown, O. 2016. Investigating the musical affordances of continuous time recurrent neural networks. In *Proceedings of the 7th International Conference on Computational Creativity (ICCC)*.
- Jaeger, H., and Haas, H. 2004. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science* 304(5667):78–80.
- Jordà, S. 2005. *Digital Lutherie: Crafting Musical Computers for New Musics' Performance and Improvisation*. Ph.D. dissertation, Universitat Pompeu Fabra.
- Kantosalo, A., and Takala, T. 2020. Five C's for human-computer co-creativity: An update on classical creativity perspectives. In *Proceedings of the 11th International Conference on Computational Creativity (ICCC)*, 17–24.
- Kiefer, C. 2014. Musical instrument mapping design with echo state networks. In *Proceedings of NIME 2014*.
- McCormack, J. 2001. Eden: An evolutionary sonic ecosystem. In *Advances in Artificial Life*, 133–142. Springer.
- Nakajima, K., and Fischer, I., eds. 2021. *Reservoir Computing: Theory, Physical Implementations, and Applications*. Springer.
- Pasquier, P.; Eigenfeldt, A.; Bown, O.; and Dubnov, S. 2016. An introduction to musical metacreation. *Computers in Entertainment* 14(2): Article 2.
- Prusinkiewicz, P., and Lindenmayer, A. 1990. *The Algorithmic Beauty of Plants*. Springer-Verlag.
- Saunders, R.; Gemeinboeck, P.; Lombard, A.; Bourke, D.; and Kocaballi, A. B. 2010. Curious Whispers: An embodied artificial creative system. In *Proceedings of the 1st International Conference on Computational Creativity (ICCC)*.
- Schaap, G., and Hedblom, M. M. 2024. Discussing the creativity of AutomaTone: An interactive music generator based on Conway's Game of Life. In *Proceedings of ICC 2024*.
- Stanley, K. O. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines* 8(2):131–162.
- Wiggins, G. A. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7):449–458.

A Implementation Details

A.1 Folding operator parameterization

Plane positions are specified by a short trigonometric sequence (four frequency and four phase parameters, plus base and step scalars), so the full sequence is reconstructible from roughly ten named floats. A random plane mode is also supported, parameterized by a single seed integer. After folding, the cloud is rotated so the final plane’s normal aligns with screen-right, giving a stable canonical orientation across configurations.

A.2 Network construction parameters

The kNN graph uses typical settings of $k = 8$ and a radius cutoff of 0.4 in normalized coordinates, constructed via a uniform spatial-hash grid. Each neuron also maintains an exponentially decaying “glow” trace used by the visualization layer. A small kickstart initially stimulates a random subset of neurons to seed activity.

A.3 Probe placement algorithm

The farthest-point sampling proceeds as follows: the first probe is the point farthest from the centroid; each subsequent probe is the argmax of the distance-to-nearest-source field after relaxing Dijkstra from all picked sources. Probe slot indices are stable across graph rebuilds, so each probe is a continuous output channel that the chart trace, the connector line, and the audio voice all key off through structural changes.

A.4 Sonifier signal chains and mappings

The oscillator path is structured as: square oscillator \rightarrow low-pass \rightarrow VCA \rightarrow panner. The granular path is: density-driven grain scheduler with position, size, and pitch jitter \rightarrow low-pass \rightarrow VCA \rightarrow panner. Both paths are summed into a shared bus passing through a compressor and reverb. Probe activation drives voice gain through a soft expander for contrast; camera-space distance modulates filter cutoff; each probe’s normalized-device-coordinate horizontal position drives stereo pan. The performer’s controls are organized into five panels (generation, visualization, network, probes, and audio) exposing seed shape, plane sequence, graph parameters, network constants, voice and pitch modes, pitch-palette size n , and granular source buffer and grain controls.

A.5 Pitch logic

Slot mode maps voice i to a pitch drawn from a maximally even set distribution (Clough and Douthett 1991) over the chromatic octave: a single integer $n \in \{1, \dots, 12\}$ selects the size of the active set, sliding voices from sparse to fully chromatic while preserving optimal even spacing.

Prune mode uses a per-point scalar computed at seed time from the azimuthal angle $\text{atan2}(z, x)/2\pi$, the same scalar that drives the body’s hue. The scalar is propagated unmodified through Ω , deduplication, and capping, so reflection is hue-invariant by construction and points related by symmetry share both color and pitch scalar in any quantization. For oscillator voices the scalar selects a pitch; for granular

voices it indexes a position in the source buffer instead, so the same geometric address drives different output spaces depending on voice path.