Spatial Frequency Discrimination of Multi-Rate Serial Signals on Untrained Neuromorphic Convolutional Hardware

Tyler Joseph Hirsch
ELLA Project | Chiang Mai, Thailand

We demonstrate that a convolutional neuromorphic processor with randomly-initialized, never-modified weights can discriminate spatial frequency patterns in binary spike frames with 100% pairwise accuracy across 15 pattern classes spanning nearly four orders of magnitude — without any training procedure. A BrainChip AKD1000 neuromorphic SoC receives 16×16×1 binary frames (spike-present: 0xFF; spike-absent: 0x00) and produces a 5-element output signature vector. The 3×3 convolutional kernel acts as a natural edge detector: uniform regions generate near-zero activation, while 0x00→0xFF transitions produce measurable output proportional to transition density. Output signatures correlate monotonically with spatial frequency (Pearson $r = -0.90$ across all output classes). Five independent random weight seeds were tested directly on AKD1000 silicon, all achieving 105/105 pairwise discriminations DISTINCT, establishing that the result is a geometric property of binary spatial convolution, not an artifact of weight selection.

The binary spike frames encode serial communication baud rates as spatial frequency patterns via the UART divisor relationship: at baud rate B sampled at crystal frequency $F = 1.8432$ MHz, each bit occupies $(D/D_{\mathrm{ref}}) \times N^2$ consecutive channels ($D = F/B$; preferred embodiment $N = 16$, 256 channels). This maps communication speed directly to spatial frequency without any temporal transform or feature extraction step. On real AKD1000 silicon, 276/276 baud-rate pairs are pairwise discriminable (L2 threshold = 0.5) at 0.6 ms per inference (~1,600 inferences/s) with fully deterministic outputs. Ten simultaneous baud rates packed into spatial regions of a single frame are classified in one inference pass; six spatial permutations of an identical rate-set produce distinct signatures (6/6), confirming that the processor preserves both spectral content and spatial arrangement. Data-content independence was confirmed across six payload sources including real UART-framed ASCII and binary data.

The approach requires no temporal windowing, no signal transformation (Gramian Angular Field, Recurrence Plot, or spectral), and no trained readout layer, distinguishing it from 1D-to-2D CNN classification methods [1, 2] and from reservoir computing approaches that train a linear classifier on random features [3, 4]. A scalar transition-counting baseline achieves 98.1% discrimination, confirming that the encoding design is the primary contribution; the neuromorphic processor closes the residual 1.9% gap on high-frequency plateau pairs where scalar methods fail. A pipeline saturation threshold $t^* = L^2/\pi^2 \approx 25.93$ (L = 16, forced by the structural coupling (ratio + 1) × pinion = 2^8) emerges from the 1D heat equation on the encoding domain [240, 256], defining the time at which pipeline feedback transients decay to e^{-1}. Validated across 11,549 hardware tests.

Table 1. Selected pairwise L2 distances on AKD1000 silicon (threshold = 0.5; N = 16).

Spatial frequency pair	L2	Result
480 vs 960 baud (single trial, seed 0)	0.000	unresolved
480 vs 960 baud (2-trial average)	1.884	DISTINCT
460,800 vs 921,600 baud (tightest single-pattern)	0.95	DISTINCT
10-baud mix: original vs reversed layout	0.59	DISTINCT
480 vs 1,843,200 baud (widest range)	30.34	DISTINCT

[1] Saxe et al., ICML Workshops, 1, 2011.
[2] Wang & Oates, AAAI Workshop, 1, 2015.
[3] Dempster et al., Data Min. Knowl. Disc., 34(5), 1454–1495, 2020.
[4] Jaeger, GMD Tech. Rep., 148(34), 13, 2001.