Radial Basis Function Networks for Spatially-Aware Predictive Coding

Hou Hei Lam

Introduction & Motivation

- The Problem with Backpropagation: While effective, backpropagation lacks biological plausibility due to its reliance on global error signals. Predictive Coding (PC) offers an alternative by using local error minimization.
- The Connectivity Gap: Standard PC implementations typically use fully connected layers. This contrasts sharply with biological neural circuits, which exhibit sparse, spatially organized "like-to-like" connectivity.
- The Hypothesis: Imposing spatial locality via Gaussian receptive fields can bridge this gap, reducing computational complexity while preserving—or even enhancing—the geometric structure of learned representations.

Biological Inspiration

- Cortical Organization: Biological cortex relies on local lateral connections to preserve the visuotopic ordering of the visual field.
- Efficiency: Structured sparsity in biological brains supports high-throughput processing. Recent mapping (MICrONS) confirms that connection strength drops with distance.
- Equivariance: Local connectivity naturally supports equivariant representations (symmetry under translation and rotation), a fundamental principle of vision.

Methodology: Radial Basis Predictive Coding

We introduce Gaussian receptive fields to restrict connectivity based on spatial proximity.

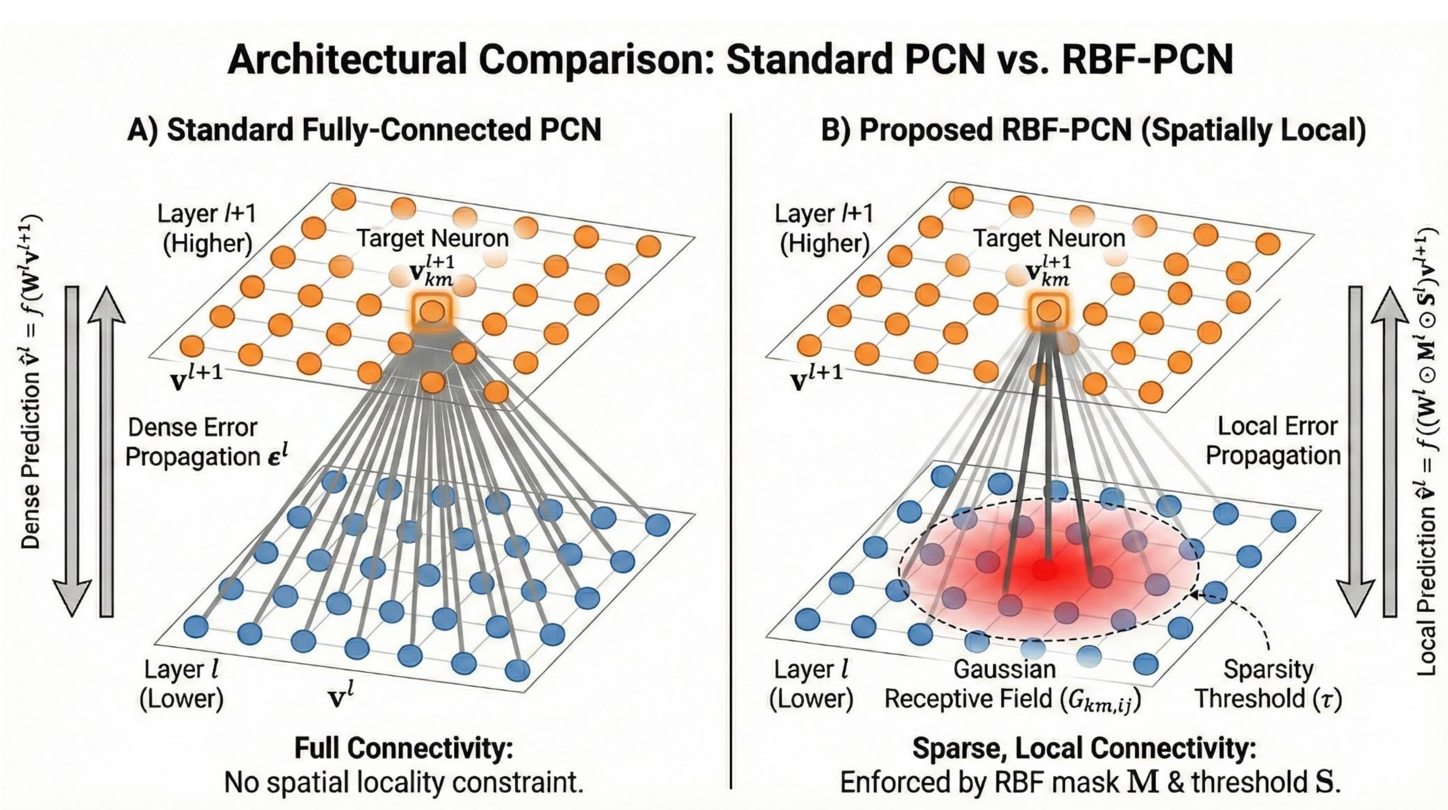

A. The RBF Constraint

For neurons on 2D grids, connectivity is defined by a Gaussian function G based on the distance between coordinates (i, j) and (k, m):

$$G_{km,ij} = exp\left(-\frac{||(k,m)-(i,j)||^2}{2\sigma^2}\right)$$

This creates a connectivity mask M applied to the weights.

B. Modified Dynamics

Top-down prediction with error $\epsilon^l = v^l - \widehat{v^l}$ becomes:

$$\widehat{v^l} = f\left((W^l \odot M^l)v^{l+1}\right)$$

Spatially constrained inference updates become:

$$\frac{dv^{l+1}}{dt} \propto \epsilon^{l+1} - \left(W^l \odot M^l\right)^T \epsilon^l$$

Complexity Reduction: Reduces complexity from $O(N_{l+1} \times N_l)$ to $O(N_{l+1} \times k)$, where k is the average connections per neuron.

C. Sparse Learning

To ensure true computational sparsity (not just weighted reduction), we employ a binary threshold mask:

- Binary Mask (S): Selects connections where the RBF value $> \tau$ (threshold).
- Effective Weights: During forward passes, weights are $(W^l \odot S^l) \odot M^l$. Weight updates only occur on active connections.

Experimental Results: Classification

We compared RBF-PCN against Standard PCN and Backpropagation (BP) on MNIST and CIFAR-10 using MLPs of varying depths (3–12 layers).

Dataset	Method	Depth 3	Depth 6	Depth 9	Depth 12
MNIST	BP	0.98 ± 0.0006	**0.98 ± 0.0017**	**0.98 ± 0.0032**	**0.97 ± 0.0049**
	PCN	**0.99 ± 0.0007**	**0.98 ± 0.0018**	0.94 ± 0.0082	0.17 ± 0.0669
	RBF-PCN	0.98 ± 0.0004	0.97 ± 0.0007	0.11 ± 0.0000	0.12 ± 0.0025
CIFAR-10	BP	**0.55 ± 0.0064**	**0.55 ± 0.0003**	**0.54 ± 0.0037**	**0.53 ± 0.0015**
	PCN	**0.55 ± 0.0027**	0.48 ± 0.0048	0.42 ± 0.0040	0.28 ± 0.0065
	RBF-PCN	0.53 ± 0.0056	0.46 ± 0.0045	0.10 ± 0.0028	0.10 ± 0.0013

Key Finding: RBF-PCN matches standard PCN in shallow networks despite sparse connectivity, but both PC variants struggle with depth.

Geometric Equivariance Analysis

We measured how well the models preserve input geometry (Translation and Rotation) by comparing feature similarity: $sim\left(f(T(x)), T(f(x))\right)$.

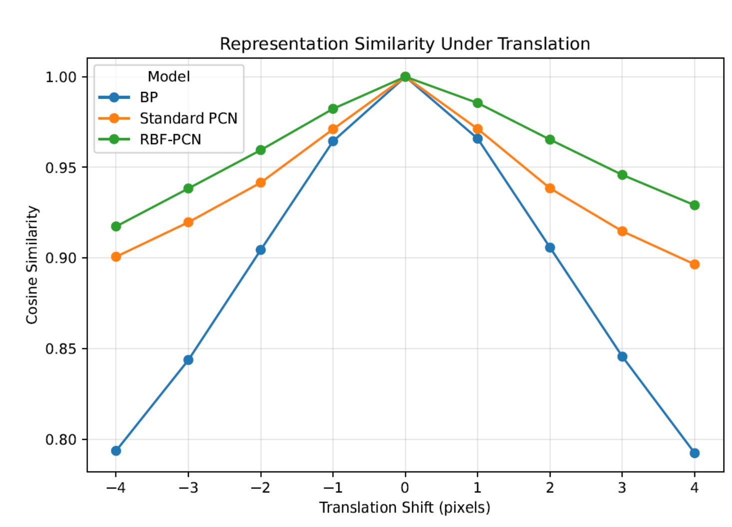

(a) CIFAR-10 translation

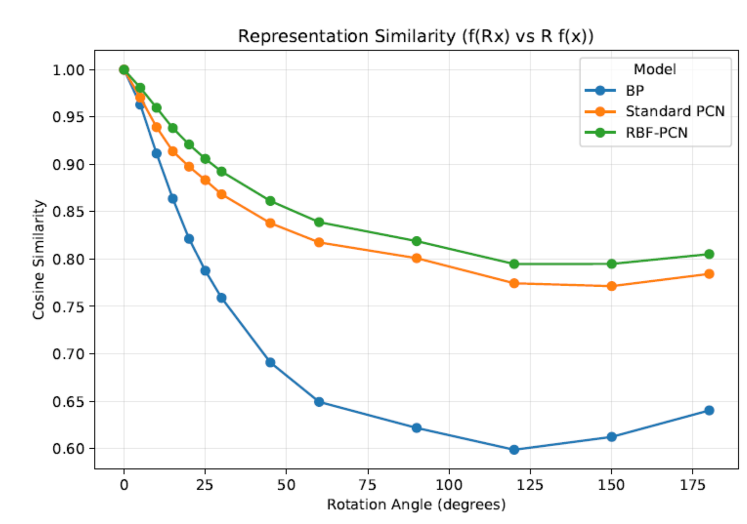

(b) CIFAR-10 rotation

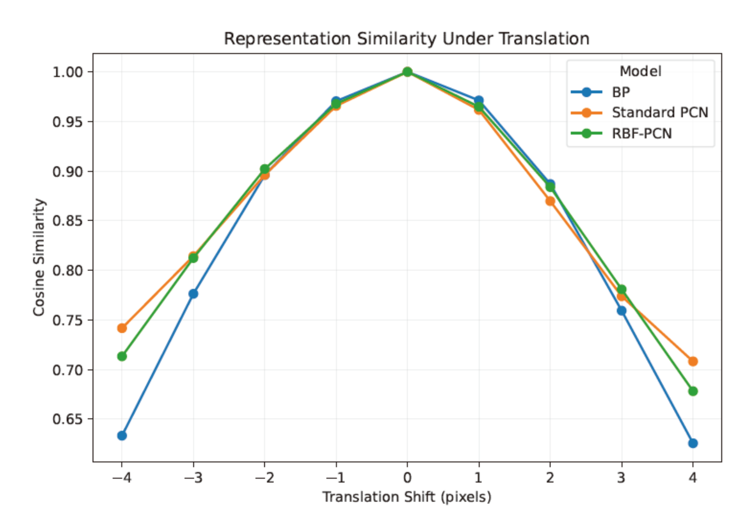

(c) MNIST translation

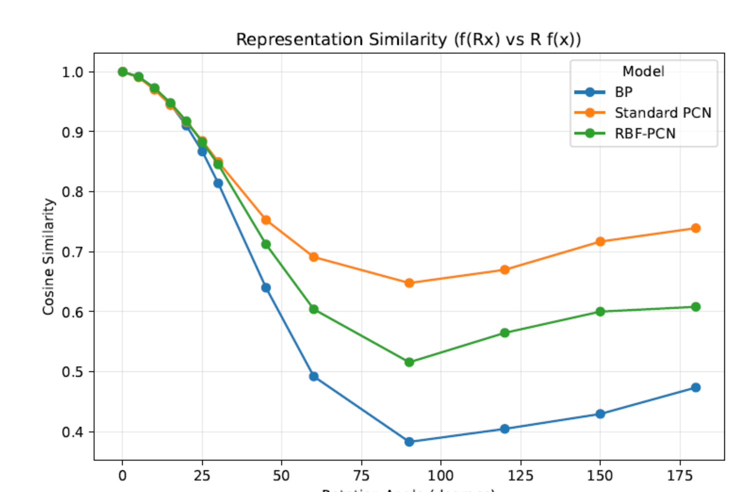

(d) MNIST rotation

- Translation: Both PCN variants (Standard and RBF) maintained higher cosine similarity under pixel shifts compared to Backpropagation.
- Rotation: RBF-PCN yielded the highest similarity on CIFAR-10, suggesting that enforcing local spatial constraints helps the network learn rotationally equivariant features better than dense connections.
- Conclusion: Predictive coding dynamics naturally favor geometric preservation over standard BP.

Geometric Equivariance Analysis

The model was tested on Image Denoising (MNIST) and Texture Synthesis (DTD dataset) to verify representational capacity.

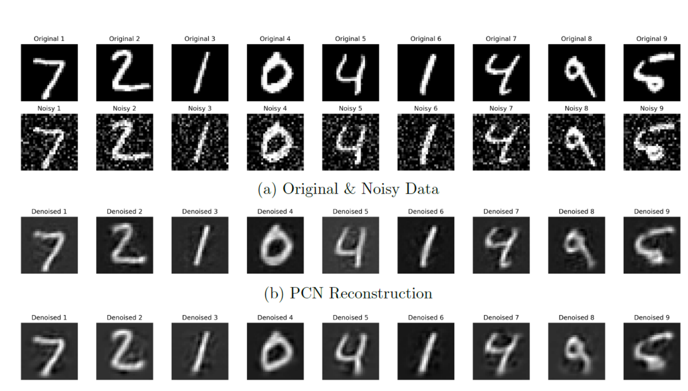

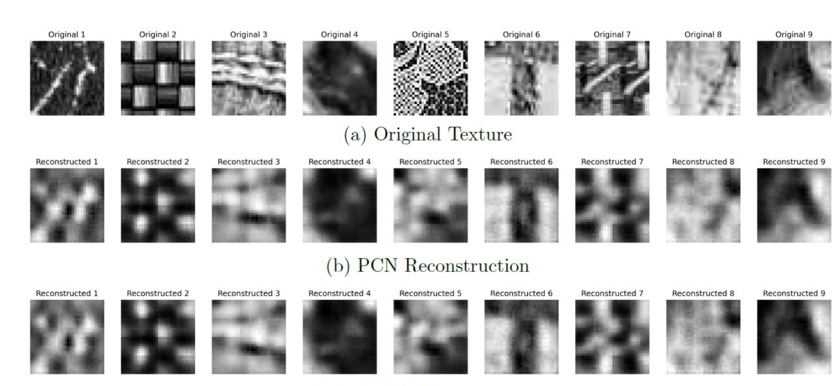

- Denoising: Depth-3 RBF-PCN achieved a PSNR of 12.84 dB, comparable to Standard PCN (13.39 dB), successfully filtering Gaussian noise.
- Texture Learning: On the DTD dataset, RBF-PCN achieved an MSE of 0.03 (identical to PCN), proving that local receptive fields are sufficient to capture complex spatial patterns.

Discussion & Future Work

RBF-PCN adds biological spatial constraints to predictive coding, matching shallow-network performance while using far fewer active parameters. Its local connectivity yields representations more robust to spatial changes than standard deep models, but performance still degrades in very deep networks, motivating work on better signal propagation (e.g., adaptive receptive fields).