Selectivity and Shape in the Design of Forward-Forward Goodness Functions

Talha Rüzgar Akkuş[*]
Poindexter Labs
tra@poindexterlabs.ai

Şuayp Talha Kocabay[*]
Poindexter Labs
stk@poindexterlabs.ai

Kamer Ali Yüksel[*]
Feedzai
kamer.yuksel@feedzai.com

Hassan Sawaf
aiXplain, Inc.
hassan@aixplain.com

Abstract

The Forward-Forward (FF) algorithm trains networks layer-by-layer using a local "goodness function," yet the choice of activity statistic remains an important design question. We systematically study how selectivity and activity shape affect FF learning, introducing selective goodness functions based on top-k and entmax-weighted energy alongside shape-sensitive objectives based on normalized higher-order moments. In particular, we propose *burstiness goodness*, a parameter-free, scale-invariant excess-kurtosis objective that rewards heavy-tailed activity patterns. Our study examines goodness functions together with activation and label-injection choices across six small-image datasets, and uses controlled experiments to isolate their effects. In validation-tuned, five-seed Fashion-MNIST experiments with matched ReLU activation, standard label injection, and per-layer scoring, burstiness improves accuracy from $73.2 \pm 0.9\%$ to $88.6 \pm 0.4\%$ over sum-of-squares. On CIFAR-10, burstiness improves block-local ResNet-18 accuracy from $28.9 \pm 0.5\%$ to $46.8 \pm 2.5\%$ under the same separate per-block label pathway. These results show that activity-shape statistics can provide effective local learning signals across the tested settings. Code is available at `https://github.com/suayptalha/ff-selectivity-shape`.

1 Introduction

The Forward-Forward (FF) algorithm [Hinton, 2022] replaces backpropagation's global backward pass with a local, layer-wise learning rule: each layer maximizes a scalar "goodness" for positive (correctly labeled) data and minimizes it for negative data. Hinton [2022] defined goodness as the sum of squared activities (SoS), and this energy-based choice has been widely used [Giampaolo et al., 2023, Lorberbom et al., 2023, Lee and Song, 2023]. A recent benchmark [Shah and Tripathi, 2025] evaluated 21 goodness functions but within a fixed architecture; no prior work has jointly studied goodness functions, activation functions, label-injection strategies, and the underlying principle governing what makes a good goodness function.

We argue this gap is significant: the goodness function defines each layer's objective landscape, determining what representations are rewarded and what features emerge. Sum-of-squares aggregates all neurons and measures total activity magnitude. In contrast, we show that shape-sensitive metrics can better capture the nuanced, non-Gaussian signatures of discriminative features. We conduct a systematic study of the following **design hypothesis**: sensitivity to the shape of neural activity can improve FF goodness objectives within the tested settings. Selective functions (top-k, entmax-weighted energy) retain magnitude information while focusing on peak activity; burstiness and normalized higher moments instead measure scale-invariant distributional shape. In matched standard-

[*]Equal contribution.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

FF experiments, burstiness improves Fashion-MNIST accuracy by 15.4 percentage points over validation-tuned SoS with ReLU held fixed. Our contribution is the formulation and controlled study of these established statistics as local FF objectives, rather than the invention of the statistics themselves.

Contributions.

1. We formulate selective top-k and entmax-weighted energy objectives and scale-invariant burstiness goodness (§4).
2. We isolate objective effects through validation-tuned, five-seed comparisons with matched activation, label pathway, and standard inference scoring (§5.2).
3. We study activation and label-pathway interactions and three parameter sweeps, and test transfer to convolutional and residual local blocks (§5.4–5.6).

2 Related Work

Forward-Forward learning. Hinton [2022] introduced FF; subsequent work studied layer sizes and negative data [Giampaolo et al., 2023], layer collaboration [Lorberbom et al., 2023], symmetric contrastive variants [Lee and Song, 2023], predictive coding [Ororbia and Mali, 2023], and per-layer label injection (FFCL; Karkehabadi et al. 2024). Shah and Tripathi [2025] benchmarked 21 goodness functions within a fixed architecture with peer normalization and downstream classifiers. Our work differs by jointly studying the goodness function, activation function, label pathway, and the scoped shape-sensitivity hypothesis.

Complete methods versus objective controls. SCFF [Chen et al., 2025] changes positive/negative construction and representation learning; Mono-Forward [Gong et al., 2026a] uses local classification losses. SymBa and FFCL likewise modify more than the activity statistic. Our loss substitutions do not benchmark these complete pipelines. Convolutional and scalable FF work [Scodellaro et al., 2025, Krutsylo, 2025, Gong et al., 2026b] addresses broader architecture and scaling questions; our spatial tests establish relative objective gains in matched local blocks, not superiority to those methods.

Sparse transformations. The α-entmax family [Martins and Astudillo, 2016, Correia et al., 2019, Peters et al., 2019] produces sparse probability distributions and has been applied to attention mechanisms. We apply entmax in a novel context: as a sparse weighting mechanism within a goodness function, creating an adaptive alternative to hard top-k selection.

Sparse coding and ICA. Sparse coding [Olshausen and Field, 1996], k-WTA networks [Ahmad and Scheinkman, 2019, Maass, 2000], and ICA [Hyvärinen and Oja, 2000] motivate our approach. Top-k encourages kWTA-like sparsity; entmax provides a differentiable relaxation. Burstiness borrows a non-Gaussianity statistic from ICA without claiming to recover independent components. The heavy-tailed distributions of deep network activations [Martin and Mahoney, 2019] further motivate kurtosis as a natural training signal. Local learning rules (Hebbian [Hebb, 1949], contrastive Hebbian [Xie and Seung, 2003], equilibrium propagation [Scellier and Bengio, 2017]) are orthogonal—we improve the objective within the FF framework.

3 Background: The Forward-Forward Algorithm

Training. Given input $\mathbf{x}$ and label y, FF creates positive/negative inputs by embedding the correct/incorrect label:

$$\mathbf{x}^+ = \text{norm}([\mathbf{x};\ s \cdot \text{onehot}(y)]), \qquad \mathbf{x}^- = \text{norm}([\mathbf{x};\ s \cdot \text{onehot}(\tilde{y})]), \quad \tilde{y} \neq y. \tag{1}$$

Each layer ℓ computes $\mathbf{h}_\ell = f_\ell(\mathbf{h}_{\ell-1}; \theta_\ell)$ and is trained via:

$$\mathcal{L}_\ell = \mathbb{E}\left[\log\left(1 + e^{\tau - g(\mathbf{h}_\ell^+)}\right)\right] + \mathbb{E}\left[\log\left(1 + e^{g(\mathbf{h}_\ell^-) - \tau}\right)\right], \tag{2}$$

where $g(\cdot)$ is the goodness function and τ a threshold. Here $\tilde{y}$ is a randomly sampled incorrect class label. Layers are trained independently; outputs are L2-normalized before propagation.

Inference. For each candidate label c, we embed and forward through all layers, predicting $\hat{y} = \arg\max_c \sum_\ell g(\mathbf{h}_\ell^{(c)})$.

The goodness function. Hinton [2022] defined $g_{\text{SoS}}(\mathbf{h}) = \sum_{i=1}^{d} h_i^2$, the energy statistic used in the original algorithm. We challenge the implicit assumption that total squared activity sufficiently summarizes a layer's representation.

4 Goodness Design: Questions and Objectives

We organize the design space around three experimental questions: does selecting peak activity help (§4.1–4.2), can scale-invariant shape statistics improve on energy (§4.3), and do these effects survive matched activations, label pathways, and architectures (§5)? The matched ReLU comparison isolates a 15.4-point gain; the spatial tests examine whether that relative objective advantage transfers beyond flattened MLPs. FFCL and activation choices below are existing components used for interaction ablations.

4.1 Top-k Goodness

Top-k goodness measures only the k most active neurons:

$$g_{\text{top-}k}(\mathbf{h}) = \frac{1}{k} \sum_{i \in \mathcal{S}_k(\mathbf{h})} h_i, \qquad \mathcal{S}_k(\mathbf{h}) = \text{argtop-}k(\mathbf{h}), \tag{3}$$

with $k = \max(5, \lfloor 0.02\, d \rfloor)$ (2% of layer width). Unlike SoS, top-k ignores the $(d-k)$ least active neurons, creating a focused learning signal that encourages sparse, discriminative representations. This design is motivated by k-winners-take-all (k-WTA) mechanisms observed in biological circuits, where competitive inhibition ensures that only the most relevant neurons respond to a given stimulus [Maass, 2000]. By optimizing for peak activity, the layer learns to allocate its capacity to only the most informative features (§6).

4.2 Entmax-Weighted Energy Goodness

Whereas top-k applies hard selection, *entmax-weighted energy* uses α-entmax [Correia et al., 2019, Peters et al., 2019] for adaptive sparse weighting:

$$g_{\text{entmax}}(\mathbf{h};\, \alpha) = \sum_{i=1}^{d} \pi_i\, h_i^2, \qquad \boldsymbol{\pi} = \text{entmax}_\alpha(\mathbf{h}). \tag{4}$$

The parameter α controls sparsity (α=1: softmax/dense; α=2: sparsemax [Martins and Astudillo, 2016]/hard sparse), and entmax *learns* how many neurons are relevant per input.

4.3 Burstiness (Excess Kurtosis) Goodness

Orthogonally to selective measurement, *burstiness goodness* measures a scale-invariant statistic of the full activation distribution—the excess kurtosis:

$$g_{\text{burst}}(\mathbf{h}) = \frac{\frac{1}{d}\sum_{i=1}^{d}(h_i - \mu)^4}{\left[\frac{1}{d}\sum_{i=1}^{d}(h_i - \mu)^2\right]^2} - 3, \tag{5}$$

where $\mu = \frac{1}{d}\sum_i h_i$, for nonzero variance. Applying LayerNorm before this statistic is mathematically redundant (up to numerical stabilization); we do not treat LN-burstiness as a separate objective. The key property is *scale invariance*: $g_{\text{burst}}(\alpha\mathbf{h}) = g_{\text{burst}}(\mathbf{h})$ for any $\alpha > 0$, making the ideal statistic invariant to uniform magnitude changes; numerical stabilization can break exact invariance. It rewards heavy-tailed ("bursty") activity patterns analogous to cortical burst firing [Lisman, 1997], connecting directly to ICA where kurtosis maximization extracts independent features [Hyvärinen and Oja, 2000]. We generalize to the p-th central moment:

$$g_{\text{moment-}p}(\mathbf{h}) = \frac{\frac{1}{d}\sum_{i=1}^{d}(h_i - \mu)^p}{\left[\frac{1}{d}\sum_{i=1}^{d}(h_i - \mu)^2\right]^{p/2}} - \beta_p, \tag{6}$$

where $\beta_p = (p-1)!!$ for even p and 0 for odd p. At $p = 4$ this recovers burstiness; higher p amplifies extreme activations.

4.4 Additional Goodness Functions

We also evaluate: *contrast top-k* ($g = \frac{1}{k}\sum_{\mathcal{S}_k^+} h_i - \frac{1}{k}\sum_{\mathcal{S}_k^-} h_i$), *LayerNorm-top-k* ($g = g_{\text{top-}k}(\text{LN}(\mathbf{h}))$), *variance* and *negative entropy*, and two external baselines from Shah and Tripathi [2025]: *softmax-energy-margin* and *game-theoretic*.

4.5 Separate Label–Feature Forwarding (FFCL)

In standard FF, labels are concatenated at the input only. FFCL [Karkehabadi et al., 2024] injects class hypotheses at *every* layer via a separate projection:

$$\mathbf{h}_\ell = \sigma(W_\ell^{\text{feat}}\,\mathbf{h}_{\ell-1}), \qquad \tilde{\mathbf{h}}_\ell = \mathbf{h}_\ell + W_\ell^{\text{label}}\,\mathbf{y}_{\text{oh}}, \tag{7}$$

where $W_\ell^{\text{label}} \in \mathbb{R}^{d\times C}$. Goodness is computed on $\tilde{\mathbf{h}}_\ell$; only the label-free $\mathbf{h}_\ell$ (L2-normalized) propagates to the next layer.

4.6 Activation Functions

We study five activations: ReLU (sparse, many zeros), GELU [Hendrycks and Gimpel, 2023] and Swish [Ramachandran et al., 2017] (smooth, dense activity), and LN-GELU / LN-Swish (pre-activation LayerNorm + smooth nonlinearity). Smooth activations help shape-sensitive goodness functions by providing richer distributions, while pre-activation normalization stabilizes inputs across layers. Algorithm 1 summarizes training.

5 Experiments

5.1 Experimental Setup

Datasets. Primary evaluation: **MNIST** [Lecun et al., 1998] and **Fashion-MNIST** [Xiao et al., 2017] (both 10-class, 28×28 grayscale; full combinatorial ablation). Cross-dataset generalization: **CIFAR-10** [Krizhevsky, 2012] (3072-dim color), **USPS** [Hull, 1994] (256-dim grayscale), **SVHN** [Netzer et al., 2011] (3072-dim color), **EMNIST-Letters** [Cohen et al., 2017] (784-dim, 26 classes). Pixel values are normalized to zero mean and unit variance.

Algorithm 1 Forward-Forward Training with Shape-Sensitive Goodness

Require: Dataset $\mathcal{D}$, goodness function g, threshold τ, label pathway (Std/FFCL)

for each layer $\ell = 1, \ldots, L$ **do**
  for each mini-batch $(\mathbf{x}, y) \sim \mathcal{D}$ **do**
    Sample wrong label $\tilde{y} \neq y$
    if Standard pathway **then**
      $\mathbf{h}_\ell^+ \leftarrow \sigma(W_\ell\,\mathbf{h}_{\ell-1}^+), \quad \mathbf{h}_\ell^- \leftarrow \sigma(W_\ell\,\mathbf{h}_{\ell-1}^-)$ // $\mathbf{h}_0^\pm$: input with label y / $\tilde{y}$
    else if FFCL pathway **then**
      $\mathbf{h}_\ell \leftarrow \sigma(W_\ell^{\text{feat}}\,\mathbf{h}_{\ell-1})$ // label-free features
      $\mathbf{h}_\ell^+ \leftarrow \mathbf{h}_\ell + W_\ell^{\text{label}}\,\text{onehot}(y), \quad \mathbf{h}_\ell^- \leftarrow \mathbf{h}_\ell + W_\ell^{\text{label}}\,\text{onehot}(\tilde{y})$
    end if
    $\mathcal{L}_\ell \leftarrow \log(1 + e^{\tau - g(\mathbf{h}_\ell^+)}) + \log(1 + e^{g(\mathbf{h}_\ell^-) - \tau})$ // g: top-k, entmax, etc.
    Update θ_ℓ via Adam on $\mathcal{L}_\ell$
  end for
  $\mathbf{h}_\ell \leftarrow \text{L2-normalize}(\mathbf{h}_\ell)$ // propagate to next layer
end for

Architecture & training. 4-layer, 2000-unit FC network (4×2000, ∼14M params). Adam [Kingma and Ba, 2017] (lr=10^{-3}), batch 500, τ=2.0, 60 epochs. Negative examples use random wrong labels. Activations are L2-normalized between layers. The combinatorial ablation uses seed 42; matched comparisons use five seeds (§5.2 and §5.6). For numerical stability, SoS is scaled by $1/d$ (mean squared activation).

Evaluation and protocol distinction. The primary controlled comparisons use only the standard sum of per-layer goodness (§3). Candidate-label inference requires one forward pass per class. The exploratory grid uses the sum of per-layer goodness and goodness of concatenated layer activations; its results are labeled accordingly. The fixed hyperparameters above describe that grid; the matched comparisons below use validation-selected settings. Std denotes input-only label concatenation; FFCL denotes the separate per-layer label pathway in Eq. 7.

Experimental grid. We evaluate 13 goodness functions crossed with 2 activations (GELU, Swish), 2 norm-gate settings[2], and 2 label pathways (standard, FFCL), plus a ReLU+SoS baseline. Additionally, we conduct three continuous sweeps: top-k cardinality k, entmax parameter α, and moment order p (§5.4).

5.2 Main Results

Does burstiness help when other design choices are held fixed? We use a held-out 55k/5k/10k train/validation/test split on Fashion-MNIST. Each objective/activation/pathway combination receives the same 3×3 validation search over label scale and threshold, followed by five test seeds. Table 1 holds architecture, activation, label pathway, and scoring fixed within each row. Standard-FF ReLU isolates a 15.4-point improvement; the per-layer-label GELU comparison isolates 5.8 points. Thus the gain survives stronger SoS tuning and does not require FFCL or concatenated-activation scoring.

Table 1: Matched Fashion-MNIST accuracy (%, mean $\pm$ standard deviation, five seeds; 4×2000; standard per-layer scoring). Each objective receives the same validation-search budget. Gains use unrounded means.

Pathway	Activation	Tuned SoS	Burstiness	Gain (pp)
Std	ReLU	73.2 ± 0.9	88.6 ± 0.4	+15.4
Std	GELU	61.5 ± 0.4	87.9 ± 0.4	+26.3
FFCL	GELU	82.1 ± 0.1	87.9 ± 0.2	+5.8

Scoring ablation. For burstiness, standard versus concatenated-ensemble scoring gives 88.6 ± 0.4 versus 88.2 ± 0.4 (Std/ReLU), 87.9 ± 0.4 versus 87.6 ± 0.4 (Std/GELU), and 87.9 ± 0.2 versus 87.8 ± 0.3 (FFCL/GELU). The extra score is therefore unnecessary for the central comparison.

5.3 Goodness Function Comparison

How do the candidate statistics compare in the exploratory grid? Table 2 summarizes the single-seed grid, selecting activation and norm-gate settings per entry. This is descriptive evidence, with untuned SoS and concatenated scoring, rather than a matched objective comparison. Burstiness and selective objectives perform well, while the FFCL interaction varies by goodness. LN-burstiness is consolidated with burstiness in the main comparison; raw numerical variants are retained only in the appendix for traceability.

[2]Norm-gating scales activations by $\sigma(\|\mathbf{h}\|) \cdot \mathbf{h}$. Across all experiments, the max accuracy difference between on/off is <0.4pp; we report the best of each pair.

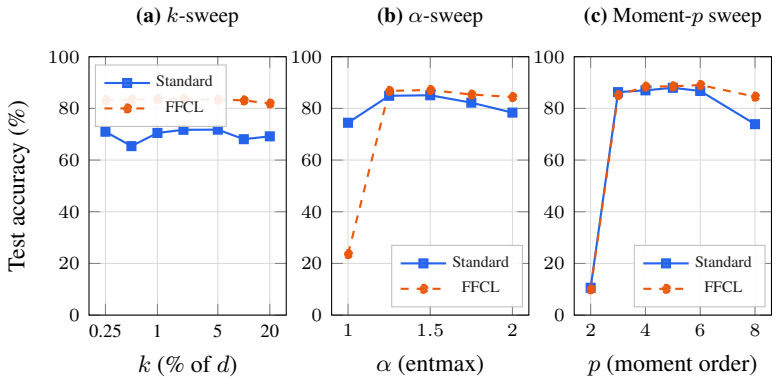

Figure 1: **Three sweep axes on Fashion-MNIST (4×2000).** The k sweep has a broad FFCL plateau; α and p have intermediate optima. **(a)** k-sweep: FFCL robust (<2pp). **(b)** α-sweep: peaks at $\alpha\approx1.5$. **(c)** Moment-p: peaks at $p\approx5$–6 (89.04% FFCL).

Table 2: Exploratory best test accuracy (%) per goodness function on Fashion-MNIST (4×2000). †From Shah and Tripathi [2025]. ‡GELU, α=1.5. §Parameter-free. Best activation/norm-gate per entry; seed 42; ensemble scoring. Std: input-only labels. Burstiness includes the equivalent LN implementation.

Goodness function	**Standard**	**FFCL**
SoS [Hinton, 2022]	61.4	82.4
Variance	61.5	81.7
Game-theoretic†	61.4	82.4
Neg. entropy	67.4	80.4
Softmax-energy-margin†	69.8	81.9
Contrast top-k	70.5	83.6
Top-k	79.0	82.9
LayerNorm-top-k	83.3	82.8
Entmax-1.5 energy‡	85.1	87.1
Burstiness (excess kurtosis)§	**88.1**	**88.4**

5.4 Sweep Analysis: Selectivity, Sparsity, and Moment Order

How much selectivity or moment sensitivity is useful? We sweep three axes on Fashion-MNIST (Figure 1): k (contrast top-k), α (entmax), and moment order p. The k sweep has a broad FFCL plateau and weak non-monotonicity under standard FF; the α and p sweeps have intermediate optima. These are sensitivity studies, not independent proof of the design hypothesis.

The k-sweep (Figure 1a) shows FFCL is remarkably robust: accuracy varies by only 1.7pp (81.9–83.6%) across a 40× range. Standard FF is more sensitive, peaking broadly around k=2–5%. The α-sweep (Figure 1b) shows a clear inverted-U peaking at $\alpha\approx1.5$ (85.1% Std, 87.1% FFCL); at α=1 (softmax), FFCL diverges entirely (23.6%), showing failure of this particular dense-weighting configuration. The moment-p sweep (Figure 1c) is the most striking: at p=2 (normalized variance), both pathways achieve only ∼10% (chance), as expected: the normalized second central moment is constant by construction. This is a sanity check, not evidence that second-order information generally fails. Performance peaks at p=5–6 (FFCL: **89.0%** at p=6) and degrades at p=8; gradient instability is a possible explanation; FFCL tolerates higher p (84.5% vs. 73.9% Std), consistent with its general robustification effect.

Shah and Tripathi [2025] reported 82.8% on Fashion-MNIST under a different pipeline with label forwarding. The held-out selection in Table 1 chose scale 10 and $\tau = 1$ for ReLU SoS. We distinguish published pipeline results from our controlled objective comparisons.

5.5 Goodness × Activation Interaction

Does the activation change the objective ranking? Table 3 reports the single-seed interaction grid. The rankings vary across activation functions. The matched ReLU and GELU comparisons in Table 1 independently evaluate the objective effect under equal validation budgets. The mechanisms behind these interactions remain to be studied.

Table 3: Exploratory goodness×activation interaction (Fashion-MNIST, 4×2000, standard FF, seed 42, ensemble scoring). Bold: best goodness per row.

	SoS	Top-k	Entmax-1.5	Burstiness
GELU	61.4	77.7	85.1	**87.2**
Swish	56.0	79.0	—	**88.1**
LN-GELU	61.5	84.0	85.4	**87.7**
LN-Swish	65.1	84.3	85.4	**88.0**

5.6 Controlled Loss Comparisons and Spatial Transfer

Does the advantage survive alternative losses and spatial blocks? In a separate controlled five-seed standard-FF protocol, replacing the loss with SymBa [Lee and Song, 2023] or SCFF [Chen et al., 2025] formulations yields Table 4. These are loss substitutions in our shared setup, not reproductions of the complete methods, their negative sampling, or their published accuracies. This protocol is reported separately from the activation/pathway decomposition in Table 1; its burstiness result is not interchangeable with that table.

Table 4: Controlled loss/statistic comparison (%, mean ± standard deviation, five seeds; standard per-layer scoring). Method names denote loss substitutions.

Loss/statistic	Fashion-MNIST	MNIST
FF / tuned SoS	73.2 ± 0.8	87.7 ± 0.5
SymBa loss / energy	73.6 ± 0.2	90.7 ± 0.1
SCFF loss / energy	73.6 ± 0.2	90.7 ± 0.1
FF / burstiness	87.3 ± 0.3	97.5 ± 0.1

For CIFAR-10, both objectives receive the same 18-configuration validation search on Conv8 with standard labels for 10 epochs. The selected hyperparameters transfer unchanged to five-seed, 30-epoch Conv8 and block-local ResNet-18 evaluations (Table 5). GELU and raw per-block goodness-sum scoring are matched; outputs are detached at local block boundaries. The double-digit improvements test architectural transfer, not the convolutional FF accuracy frontier.

Table 5: CIFAR-10 spatial transfer (%, mean ± standard deviation, five seeds). Std: input-only labels; separate: per-block label pathway.

Architecture	Labels	SoS	Burstiness	Gain (pp)
Conv8	Std	27.7 ± 0.6	40.6 ± 1.8	+12.9
Conv8	Separate	25.0 ± 0.7	46.2 ± 0.7	+21.2
ResNet-18	Separate	28.9 ± 0.5	46.8 ± 2.5	+17.8

5.7 Label-Pathway Interaction

How does label forwarding interact with goodness? In the matched GELU comparisons in Table 1, burstiness outperforms SoS under both standard input-only labels and separate per-layer label forwarding. The size of the gap depends on the pathway, so the effect of label injection and the goodness function should be assessed together.

5.8 Cross-Dataset Generalization

How do the objectives behave across datasets? Table 6 reports single-seed results across six datasets under the exploratory ensemble-scoring protocol. The listed configurations differ in activation and label pathway, and therefore do not isolate goodness effects. Burstiness configurations achieve strong accuracy across the tested datasets; matched comparisons in Tables 1 and 5 provide the controlled tests of objective effects.

Table 6: Exploratory test accuracy (%) across six datasets (4×2000, seed 42, ensemble scoring). Rows use different activations and label pathways.

	MNIST	F-MNIST	CIFAR-10	USPS	SVHN	EMNIST
FFCL + SoS (GELU)	93.5	82.4	46.5	76.8	71.7	77.8
Std + Top-k (Swish)	90.2	78.6	37.4	82.9	57.7	66.3
FFCL + Ctrst. top-k	93.3	83.6	47.1	90.0	66.0	79.7
Std + Burstiness (Swish)	96.5	88.1	53.6	93.1	78.8	86.2
FFCL + Burstiness (GELU)	98.1	88.4	54.4	93.3	82.8	**91.6**
Std + Burstiness (LN-GELU)	96.9	87.7	53.3	**93.9**	81.7	89.0
FFCL + Burstiness (LN-GELU)	**98.2**	**88.7**	54.8	93.4	**84.0**	91.2
Std + Burstiness (LN-Swish)	97.2	88.0	54.3	92.7	82.1	89.1
FFCL + Burstiness (LN-Swish)	98.1	88.7	**55.6**	92.6	83.7	91.1

6 Analysis

Our matched results support a scoped design hypothesis: activity-shape statistics can provide useful local learning signals beyond energy in the tested FF settings. This is an existence result, not a necessary or sufficient condition for effective goodness, nor a universal law across architectures and datasets. Selective statistics still retain magnitude information, whereas kurtosis removes uniform scale.

The ICA connection is motivational: kurtosis is a classical non-Gaussianity criterion in ICA [Hyvärinen and Oja, 2000], but maximizing coordinate-wise kurtosis in supervised FF does not establish independent components or perform a full ICA algorithm. Likewise, intermediate sweep optima can arise from ordinary hyperparameter sensitivity. The constant $p = 2$ statistic is only a sanity check. A falsifiable future diagnostic is an orthogonal rotation of the activation coordinates: it preserves SoS energy while generally changing coordinate-wise kurtosis. We have not run that diagnostic or a dedicated shape-sensitive negative control.

Activation interaction. The exploratory grid shows that the preferred activation depends on the goodness statistic; the tuned comparison also finds a weaker SoS result with GELU than ReLU. ReLU sparsity and smooth-activation density offer a possible interpretation, but we do not directly measure the proposed signal-dilution mechanism. In particular, shape-sensitive objectives are not guaranteed to eliminate background noise.

7 Discussion and Limitations

Accuracy and scope. Matched goodness-only comparisons provide the primary evidence. Published backpropagation and FF accuracies use different pipelines and cannot be directly compared without matching their training and evaluation protocols. The six-dataset grid remains exploratory and mostly single-seed; only subsets receive repeated evaluation. Spatial transfer is limited to CIFAR-10 local Conv8 and ResNet-18; larger-scale generalization remains open.

Reproducibility and cost. The matched MLP and spatial tests use five seeds; the broad combinatorial grid uses seed 42 and concatenated-ensemble scoring. Standard per-layer scoring is primary. Candidate-label inference still requires a forward pass for each class, limiting scalability as class count grows. In the spatial experiments burstiness increases inference time by approximately 16–21% and training time by 24–29% over SoS under the same candidate-label protocol; these are setting-specific costs. The statistic has no shape hyperparameter, but the training protocol still requires tuning.

Hyperparameters. Top-k and entmax each introduce one hyperparameter, but our sweeps show FFCL configurations are remarkably robust ($<$2pp variation for k, $<$3pp for α). Burstiness is entirely parameter-free. The moment-p family has $p \in [4, 6]$ as a broad, stable optimum.

8 Conclusion

The choice of goodness function substantially affects Forward-Forward learning. We studied selective objectives that focus on peak activities and shape-sensitive objectives based on normalized moments, including burstiness, a parameter-free and scale-invariant excess-kurtosis statistic. Our experiments show that these choices interact with activation functions and label pathways, motivating comparisons that hold these components fixed when assessing an objective.

In validation-tuned, five-seed Fashion-MNIST experiments, burstiness improves accuracy from $73.2 \pm 0.9\%$ to $88.6 \pm 0.4\%$ over SoS with matched ReLU activations, standard label injection, and per-layer scoring. The advantage also persists in controlled comparisons with separate per-layer labels and in local convolutional and residual blocks on CIFAR-10. Together, these findings show that activity-shape statistics can provide effective local learning signals across the tested settings and architectures.

The broader goodness-function grid and parameter sweeps reveal useful variation across objectives, while also highlighting sensitivity to experimental choices. Future work can test shape-sensitive objectives on larger datasets and architectures, examine their behavior under changes of activation coordinates, and study their interaction with alternative local learning procedures and inference strategies.

Acknowledgements and Disclosures

The authors received no third-party funding or support for the research reported in this paper. We thank Poindexter Labs for its support for in-person attendance at the conference. The authors declare no other competing interests.

References

Subutai Ahmad and Luiz Scheinkman. How can we be so dense? the benefits of using highly sparse representations, 2019. URL https://arxiv.org/abs/1903.11257.

Xing Chen, Dongshu Liu, Jérémie Laydevant, and Julie Grollier. Self-contrastive forward-forward algorithm. *Nature Communications*, 16(1):5978, Jul 2025. ISSN 2041-1723. doi: 10.1038/s41467-025-61037-0. URL https://doi.org/10.1038/s41467-025-61037-0.

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 2921–2926, 2017. doi: 10.1109/IJCNN.2017.7966217.

Gonçalo M. Correia, Vlad Niculae, and André F. T. Martins. Adaptively sparse transformers, 2019. URL https://arxiv.org/abs/1909.00015.

Fabio Giampaolo, Stefano Izzo, Edoardo Prezioso, and Francesco Piccialli. Investigating random variations of the forward-forward algorithm for training neural networks. In *2023 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7, 2023. doi: 10.1109/IJCNN54540.2023.10191727.

James Gong, Bruce Li, and Waleed Abdulla. Mono-forward: Revisiting forward-forward through objective-locality decomposition, 2026a. URL https://arxiv.org/abs/2501.09238.

Qingchun Gong, Robert Bogdan Staszewski, and Kai Xu. Adaptive spatial goodness encoding: Advancing and scaling forward-forward learning without backpropagation, 2026b. URL https://arxiv.org/abs/2509.12394.

D. O. Hebb. *The organization of behavior; a neuropsychological theory.* The organization of behavior; a neuropsychological theory. Wiley, Oxford, England, 1949.

Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus), 2023. URL https://arxiv.org/abs/1606.08415.

Geoffrey Hinton. The forward-forward algorithm: Some preliminary investigations, 2022. URL https://arxiv.org/abs/2212.13345.

J.J. Hull. A database for handwritten text recognition research. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(5):550–554, 1994. doi: 10.1109/34.291440.

A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4):411–430, 2000. ISSN 0893-6080. doi: https://doi.org/10.1016/S0893-6080(00)00026-5. URL https://www.sciencedirect.com/science/article/pii/S0893608000000265.

Ali Karkehabadi, Houman Homayoun, and Avesta Sasan. Ffcl: Forward-forward net with cortical loops, training and inference on edge without backpropagation, 2024. URL https://arxiv.org/abs/2405.12443.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL https://arxiv.org/abs/1412.6980.

Alex Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 05 2012.

Andrii Krutsylo. Scalable forward-forward algorithm, 2025. URL https://arxiv.org/abs/2501.03176.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Heung-Chang Lee and Jeonggeun Song. Symba: Symmetric backpropagation-free contrastive learning with forward-forward algorithm for optimizing convergence, 2023. URL https://arxiv.org/abs/2303.08418.

J. E. Lisman. Bursts as a unit of neural information: making unreliable synapses reliable. *Trends in Neurosciences*, 20(1):38–43, jan 1997. doi: 10.1016/S0166-2236(96)10070-9.

Guy Lorberbom, Itai Gat, Yossi Adi, Alex Schwing, and Tamir Hazan. Layer collaboration in the forward-forward algorithm, 2023. URL https://arxiv.org/abs/2305.12393.

W. Maass. On the computational power of winner-take-all. *Neural Computation*, 12(11):2519–2535, nov 2000. doi: 10.1162/089976600300014827.

Charles H. Martin and Michael W. Mahoney. Traditional and heavy-tailed self regularization in neural network models, 2019. URL https://arxiv.org/abs/1901.08276.

André F. T. Martins and Ramón Fernandez Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification, 2016. URL https://arxiv.org/abs/1602.02068.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Ng. Reading digits in natural images with unsupervised feature learning. *NIPS*, 01 2011.

Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, jun 1996. ISSN 1476-4687. doi: 10.1038/381607a0. URL https://doi.org/10.1038/381607a0.

Alexander Ororbia and Ankur Mali. The predictive forward-forward algorithm, 2023. URL https://arxiv.org/abs/2301.01452.

Ben Peters, Vlad Niculae, and André F. T. Martins. Sparse sequence-to-sequence models, 2019. URL https://arxiv.org/abs/1905.05702.

Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions, 2017. URL https://arxiv.org/abs/1710.05941.

Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation, 2017. URL https://arxiv.org/abs/1602.05179.

Riccardo Scodellaro, Ajinkya Kulkarni, Frauke Alves, and Matthias Schröter. Training convolutional neural networks with the forward–forward algorithm. *Scientific Reports*, 15(1), November 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-26235-2. URL http://dx.doi.org/10.1038/s41598-025-26235-2.

Arya Shah and Vaibhav Tripathi. In search of goodness: Large scale benchmarking of goodness functions for the forward-forward algorithm, 2025. URL https://arxiv.org/abs/2511.18567.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017. URL https://arxiv.org/abs/1708.07747.

Xiaohui Xie and H. Sebastian Seung. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural Computation*, 15(2):441–454, feb 2003. doi: 10.1162/089976603762552988.

A Experimental Details

Label embedding. Following Hinton [2022], input images are flattened to a 784-dimensional vector and concatenated with a one-hot label vector scaled by $s = 5.0$, yielding a 794-dimensional input that is L2-normalized. For FFCL, the input to the first layer uses only the 784-dimensional image (no label concatenation), and labels are injected at every layer via the per-layer label projection (Eq. 7).

Top-k parameter. For the 4×2000 architecture, $k = \max(5, \lfloor 0.02 \times 2000 \rfloor) = 40$ neurons (2% of layer width). For contrast top-k, $k = \max(5, \lfloor 0.01 \times 2000 \rfloor) = 20$ neurons (1%). The sparsity sweep in §5.4 varies k from 0.25% to 20% (5 to 400 neurons).

Norm-gating. Norm-gating applies $\sigma(\|\mathbf{h}\|) \cdot \mathbf{h}$ after activation, where σ is the sigmoid function. This rescales the activation vector based on its norm before passing to the next layer. In all experiments, the maximum accuracy difference between norm-gate on/off (same goodness/activation) was 0.39pp; we report the best of each pair in the main text.

Entmax computation. We use the `entmax` package [Correia et al., 2019] which implements α-entmax via the bisection algorithm. For the sparsity sweep, α ranges from 1.0 (softmax) to 2.0 (sparsemax) in steps of 0.25. Entmax-weighted energy (§4.2) first applies α-entmax to the activation vector to obtain sparse weights, then computes the weighted sum of squared activations.

External baseline implementations. *Softmax-energy-margin* [Shah and Tripathi, 2025]: We implement this as $g = g_{\text{SoS}} + 0.5 \cdot \bar{g}_{\text{SoS}} \cdot (-\log \sum_i \exp(h_i/T))$ with temperature $T = 1.0$ and margin $\lambda = 0.5$, where $\bar{g}_{\text{SoS}}$ is a running mean of goodness values. *Game-theoretic*: We implement this as a weighted SoS where each neuron's squared activation is weighted by its relative magnitude $w_i = |h_i|/(\sum_j |h_j| + \epsilon)$.

Compute resources. Experiments were run on NVIDIA A100 GPUs. Each standard FF experiment takes 30–60 seconds. Each FFCL experiment takes 35–70 seconds (the label projection adds minimal overhead). Each entmax experiment takes 200–400 seconds due to the bisection-based entmax computation. The exploratory suite (including all goodness functions, both label pathways, the sparsity sweep, and both datasets) completes in approximately 4 GPU-hours.

B MNIST Results

Table 7 presents the complete MNIST results (4×2000). On this task, the best configurations exceed 88% accuracy, and the differences between goodness functions are smaller than on Fashion-MNIST. FFCL provides a consistent ∼3–5pp lift.

Table 7: Test accuracy (%) on MNIST (4×2000), ranked by accuracy. Best of norm-gate on/off reported for each configuration. Highlighted : burstiness variants.

#	Label	Act.	Goodness	Acc%
1	FFCL	GELU	burstiness	**98.09**
2	FFCL	GELU	LN-burstiness	98.09
3	FFCL	Swish	burstiness	98.00
4	FFCL	Swish	LN-burstiness	98.00
5	Std	Swish	burstiness	96.54
6	Std	Swish	LN-burstiness	96.45
7	Std	GELU	LN-burstiness	93.64
8	FFCL	GELU	SoS	93.45
9	FFCL	GELU	game-theoretic	93.45
10	FFCL	GELU	contrast top-k	93.34
11	FFCL	GELU	variance	93.22
12	FFCL	Swish	game-theoretic	93.22
13	FFCL	Swish	SoS	93.21
14	FFCL	Swish	contrast top-k	93.14
15	FFCL	GELU	top-k	92.99
16	FFCL	Swish	top-k	92.93
17	Std	GELU	burstiness	92.49
18	FFCL	Swish	variance	92.23
19	Std	GELU	LN-top-k	92.06
20	FFCL	GELU	softmax-e-m	91.61
21	Std	Swish	LN-top-k	90.80
22	FFCL	Swish	softmax-e-m	90.40
23	Std	GELU	top-k	90.37
24	Std	Swish	top-k	90.23
25	FFCL	GELU	LN-top-k	89.34
26	FFCL	Swish	LN-top-k	89.08
27	Std	ReLU	SoS (baseline)	88.76
28	FFCL	Swish	entropy	88.50
29	FFCL	GELU	entropy	88.17
30	Std	GELU	contrast top-k	86.90
31	Std	Swish	contrast top-k	86.45
32	Std	Swish	SoS	82.54
33	Std	Swish	game-theoretic	82.51
34	Std	GELU	game-theoretic	78.06
35	Std	GELU	SoS	77.93

In this single-seed exploratory grid, FFCL + GELU + burstiness achieves 98.09% on MNIST. The validation-tuned, five-seed MNIST comparison is reported in Table 4.

C Complete Fashion-MNIST Results

Table 8 presents the complete ranked results for all configurations on Fashion-MNIST (4×2000), including both standard and FFCL label pathways.

Table 8: Complete ranked results on Fashion-MNIST (4×2000). Best of norm-gate on/off reported. Highlighted : burstiness variants.

#	Label	Act.	Goodness	Acc%
1	FFCL	GELU	burstiness	**88.41**
2	FFCL	GELU	LN-burstiness	88.41
3	FFCL	Swish	burstiness	88.21
4	FFCL	Swish	LN-burstiness	88.21
5	Std	Swish	LN-burstiness	88.11
6	Std	Swish	burstiness	88.07
7	Std	GELU	burstiness	87.15
8	Std	GELU	LN-burstiness	87.14
9	FFCL	GELU	contrast top-k	83.59
10	FFCL	Swish	contrast top-k	83.55
11	Std	GELU	LN-top-k	83.28
12	FFCL	Swish	top-k	82.93
13	FFCL	Swish	LN-top-k	82.75
14	Std	Swish	LN-top-k	82.71
15	FFCL	GELU	top-k	82.65
16	FFCL	GELU	game-theoretic	82.38
17	FFCL	GELU	SoS	82.38
18	FFCL	GELU	LN-top-k	82.25
19	FFCL	Swish	game-theoretic	82.18
20	FFCL	Swish	SoS	82.17
21	FFCL	GELU	softmax-e-m	81.89
22	FFCL	GELU	variance	81.74
23	FFCL	Swish	variance	81.29
24	FFCL	Swish	softmax-e-m	80.98
25	FFCL	Swish	entropy	80.43
26	FFCL	GELU	entropy	79.93
27	Std	Swish	top-k	79.03
28	Std	GELU	top-k	77.65
29	Std	GELU	contrast top-k	70.49
30	Std	Swish	softmax-e-m	69.85
31	Std	GELU	softmax-e-m	68.72
32	Std	Swish	contrast top-k	67.99
33	Std	Swish	entropy	67.39
34	Std	GELU	entropy	66.07
35	Std	GELU	variance	61.55
36	Std	GELU	SoS	61.43
37	Std	GELU	game-theoretic	61.37
38	Std	Swish	variance	60.72
40	Std	Swish	SoS	55.99
41	Std	Swish	game-theoretic	55.99

Burstiness and LN-burstiness occupy all top-8 positions. This table does *not* include entmax or moment-p results, which are presented in the sweep (Table 9). When sweep results are included, FFCL + moment-$p=6$ (89.04%) takes the overall top position, with FFCL + moment-$p=5$ (88.55%) and FFCL + burstiness (88.41%) closely following.

D Sparsity Sweep: Full Results

Table 9 presents the complete numerical results from the sparsity sweep (Figure 1).

Table 9: Complete sweep results on Fashion-MNIST (4×2000, GELU, no norm-gate).

(a) k-sweep

k (%)	Std	FFCL
0.25	70.97	83.13
0.5	65.40	83.42
1	70.49	83.59
2	**71.71**	**83.63**
5	71.77	83.42
10	68.08	83.10
20	69.18	81.89
Range	6.37	1.74

(b) α-sweep

α	Std	FFCL
1.00	74.44	23.60*
1.25	84.87	86.77
1.50	**85.08**	**87.12**
1.75	82.21	85.36
2.00	78.36	84.41
Range†	6.72	2.71

*Diverged. †Excl. $\alpha = 1$.

(c) Moment-p sweep

p	Std	FFCL
2	10.57*	10.00*
3	86.29	85.15
4	86.98	88.41
5	87.96	88.55
6	86.74	**89.04**
8	73.89	84.54
Range†	14.07	4.50

*Constant-statistic sanity check. †Excl. $p = 2$.

Key observations. (1) FFCL is dramatically more robust than standard FF across all three sweep axes: FFCL ranges are 1.74pp (k), 2.71pp (α), and 4.50pp (p, excl. $p = 2$), compared to 6.37pp, 6.72pp, and 14.07pp for standard. (2) The k sweep has a broad FFCL plateau; α and p have intermediate optima. (3) The moment-p sweep reveals the highest absolute accuracies (89.04% at $p = 6$ for FFCL), showing the strongest result within this exploratory sweep. (4) At $p = 2$ the normalized statistic is constant by construction, so chance accuracy is a sanity check.

E Burstiness: Detailed Results

Table 10 presents a detailed breakdown of burstiness results across all activation and label-pathway combinations on Fashion-MNIST (4×2000).

Table 10: Burstiness variants on Fashion-MNIST (4×2000, best of norm-gate on/off). LN-burstiness is an equivalent numerical implementation, not a distinct objective. Raw runs are retained here for traceability. LN-GELU and LN-Swish activations improve over GELU/Swish with FFCL.

Label	Act.	Goodness	Acc%
FFCL	LN-GELU	LN-burstiness	**88.71**
FFCL	LN-GELU	burstiness	88.69
FFCL	LN-Swish	burstiness	88.68
FFCL	LN-Swish	LN-burstiness	88.66
FFCL	GELU	burstiness	88.41
FFCL	GELU	LN-burstiness	88.41
FFCL	Swish	burstiness	88.21
FFCL	Swish	LN-burstiness	88.21
Std	Swish	LN-burstiness	88.11
Std	Swish	burstiness	88.07
Std	LN-Swish	burstiness	87.98
Std	LN-Swish	LN-burstiness	87.98
Std	LN-GELU	LN-burstiness	87.67
Std	LN-GELU	burstiness	87.66
Std	GELU	burstiness	87.15
Std	GELU	LN-burstiness	87.14

Key observations. (1) LayerNorm before kurtosis is mathematically redundant for nonzero variance; small numerical differences in these runs are not evidence for a separate method. (2) Pre-activation normalization (LN-GELU, LN-Swish) consistently outperforms plain GELU/Swish with FFCL: LN-GELU reaches 88.71% and LN-Swish 88.68%, vs. 88.41% for GELU and 88.21% for Swish. For standard FF, LN-Swish (87.98%) slightly trails Swish (88.07%) but beats LN-GELU (87.66%) and GELU (87.15%). (3) The label-pathway interaction varies by activation in this single-seed grid.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction foreground matched, validation-tuned, five-seed goodness comparisons and standard per-layer scoring (Table 1), with spatial transfer in Table 5. Shape sensitivity is a design hypothesis within the tested settings; exploratory configuration gains are distinguished from isolated objective effects.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 7 discusses the absolute accuracy gap, single seed, dataset scope, computational cost of entmax, and hyperparameter sensitivity.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [N/A]

Justification: The paper provides informal analysis (Section 6) rather than formal theorems.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Experimental methods and settings are described in Sections 4–5.1 and Appendix A; additional controlled-experiment configurations and code are available in the linked repository.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code is linked in the abstract, including experiment scripts and machine-readable results. MNIST, Fashion-MNIST, CIFAR-10, USPS, SVHN, and EMNIST are publicly available standard benchmarks.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 5.1 and Appendix A describe the experimental settings; additional controlled-experiment configurations are provided in the linked repository.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The combinatorial ablation uses a single seed (42). The central matched comparisons and spatial tests use five seeds (Tables 1 and 5); error bars denote standard deviations.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendix A specifies that experiments were run on NVIDIA A100 GPUs, with per-experiment timing (30–60s for standard FF, 200–400s for entmax) and exploratory compute ($\sim$4 GPU-hours). Spatial relative timings are reported in Section 7; aggregate compute for the controlled experiments is not available in the manuscript.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: This is a fundamental research contribution on local learning rules with no direct negative societal impact.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

Justification: This is a fundamental study of local learning rule design. Potential applications (e.g., neuromorphic hardware, on-device learning) are speculative at this stage.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The work involves standard benchmark datasets (MNIST, Fashion-MNIST, CIFAR-10, USPS, SVHN, EMNIST) and small models with no safety concerns.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All six datasets used (MNIST, Fashion-MNIST, CIFAR-10, USPS, SVHN, EMNIST) are public standard benchmarks and are properly cited. The PyTorch framework and the `entmax` package, which form the basis of our supplementary code, are used under their respective open-source licenses (BSD and MIT).

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The supplementary code includes a README with instructions, a requirements file, and comments explaining all goodness function implementations.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: No human subjects research was conducted.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification: No human subjects research was conducted.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs were used exclusively for writing and editing purposes, which does not impact the core methodology, scientific rigor, or originality of this research.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.