arXiv:2611.31513v1 [cs.LG] 23 Nov 2026

How to build and train your recurrent-depth model: A fixed-point perspective

Thomas Pethick
tmpethick@gmail.com

Abstract

Recurrent-depth models are trained with truncated backpropagation, yet it remains unclear why this works. We show that input injection, commonly used in recurrent-depth models, is exactly Halpern iteration, and that a slowly vanishing stepsize provably converges with a backpropagation window polylogarithmic in the depth. The guarantees require a Lipschitz operator, which for SwiGLU leads to a weight constraint in both row and spectral norm, while only a row norm constraint for gates. Based on these constraints, we develop an optimizer tailored to each parameter group, Sparrow, which provides a drop-in replacement of Muon/MuonH with no overhead, that speeds up training even beyond recurrent models.

github.com/tmpethick/sparrow

1 Introduction

Recurrent-depth models such as the latent reasoning model Huginn-0125 (Geiping et al., 2025) scale test-time computation by repeatedly applying a shared block to a latent state. Remarkably, they are trained with *truncated* backpropagation, differentiating through only the last few iterations of the recurrence.

At first sight it seems almost magical that this works. The truncated gradient ignores how all earlier iterations depend on the parameters, and nothing a priori prevents the latent state from drifting as the number of iterations grows. One component that appears to be important in these architectures is *input injection*, which feeds the encoded query back into every iteration (Bansal et al., 2022; Geiping et al., 2025).

In a seemingly separate vein, constraining or normalizing the *rows* of weight matrices has recently become widespread in both architectures and optimizers such as nGPT, NorMuon, Aurora and Muown (Loshchilov et al., 2025; Li et al., 2025; Dewulf et al., 2026; Lion et al., 2026), often in combination with the spectral updates of Muon (Jordan et al., 2024). It is less clear when such row constraints help, and on which matrices they should be imposed.

As we will see, these two threads are closely connected, a connection that will be revealed by asking the following question:

How should we structure the forward pass for truncated backpropagation to provably converge?
How should we train the resulting model?

To reason about these questions, we treat the recurrence as a fixed-point problem. The model takes in some encoded query u and iterates a learnable refinement map, the forward operator $\mathbb{T}_\theta$, multiple times. We can extract an unreasonable amount from this perspective.

First, it turns out that additive input injection is exactly an instance of Halpern iteration (Halpern, 1967) with a fixed stepsize. The anchoring to the query makes each iteration a contraction even when $\mathbb{T}_\theta$ is mildly expansive, which provides a possible explanation for why truncated backpropagation works. However, it also reveals that the iterates do *not* converge to a fixed point of $\mathbb{T}_\theta$, but to a biased solution. We will see that this bias can be mitigated by scheduling the stepsize carefully.

Second, all of these guarantees hinge on controlling the Lipschitz constant of $\mathbb{T}_\theta$. In particular, for SwiGLU blocks (Shazeer, 2020), it turns out that the gating makes the Lipschitz constant depend on a product of row and spectral norms, so the gate and value matrices should be constrained to the *intersection* of the spectral and row-norm balls, which motivates a new optimizer. It might sound like a small subset of matrices, but the value and gate matrices alone makes up 44% of the non-embedding model parameters in a standard decoder-only Transformer. These conclusions are not restricted to recurrent models, but inform how to train general deep models even without weight-tying.

Contributions Concretely, we make the following contributions:

(i) *Halpern iteration for recurrent depth.* We formulate a set of reasonable desiderata for a recurrent forward pass, namely existence of (potentially multiple) fixed points with input-dependent selection and a truncated backpropagation window growing only polylogarithmically with the depth (Desiderata 1 to 3). Halpern iteration, which we show is equivalent to input injection, falls out as one way of satisfying all three desiderata with a carefully scheduled anchor weight. Moreover, we show that Halpern iteration enjoys a robustness guarantee.

(ii) *A Lipschitz neural parameterization.* Since these results rely on the Lipschitz condition of the neural parameterization $\mathbb{T}_\theta$, we derive Lipschitz certificates for common blocks such as ReLU^2 based MLPs, SwiGLU, gated attention, and gated SSMs in Section 5. This perspective reveals that SwiGLU and squared ReLU require a joint row-norm and spectral-norm constraint, while gates only require a row-norm constraint. As a side-effect we also find that pre-norm is essential specifically for SwiGLU/ReLU^2 and attention, and that normalization layers in general ensure existence of fixed points.

(iii) *Sparrow: a drop-in replacement of Muon.* We propose Sparrow, a parameter-group dependent optimizer that is tailored to these constraints (Section 6). It applies the spectral oracle at most once, so it has the same cost per step as Muon. Moreover, we derive a predetermined choice of constraint radii, so that no retuning is required.

(iv) *Experiments.* We show that the joint constraint is simultaneously more expressive and more robust than a spectral ball for SwiGLU blocks, and characterize when the two coincide (Section 7.1). Without any retuning or additional computational overhead, Sparrow speeds up training of a small language model over Muon both with and without looping (Section 7.2). We find that when combining Sparrow with weight normalization as in MuonH one can ignore normalizing the update direction.

2 Related work

Fixed-point iterations The two classical schemes for finding fixed points date back to Krasnosel'skiĭ–Mann (KM) iteration (Mann, 1953), which relaxes toward the current operator output, and Halpern iteration (Halpern, 1967), which repeatedly injects a fixed anchor. The latter has received a lot of attention in recent years, due to the schedule $\beta_k = 1/(k+2)$ achieving an accelerated rate (Sabach and Shtern, 2017; Lieder, 2021), which is worst-case optimal (Park and Ryu, 2022). Our work departs from this choice of stepsize and focuses on showing compatibility with truncated backpropagation with constant and logarithmic stepsize choices. The closest work we are aware of is Diakonikolas (2025) who also study Halpern with fixed stepsize.

Classically, analysis of fixed point iterations relies on nonexpansiveness of the operator, see e.g, Bauschke and Combettes (2017) for a standard treatment. In recent years the community has moved beyond nonexpansiveness with structures such as conic (quasi-)nonexpansiveness (Bauschke et al., 2021; Alacaoglu et al., 2024; Pethick et al., 2023, 2025b). We instead follow Diakonikolas (2025) by allowing for mildly expansive operators in the case of fixed stepsize Halpern.

Implicit and recurrent-depth models Deep equilibrium models (Bai et al., 2019; Winston and Kolter, 2020; Sato and Iiduka, 2026) define the output as a fixed point of a weight-tied layer and compute gradients by implicit differentiation at that fixed point, which requires solving a linear system in the backward pass. Recurrent-depth models (Bansal et al., 2022; Geiping et al., 2025) instead unroll the iteration explicitly and train with truncated backpropagation through the last iterations. We directly adopt the input-injection structure of Geiping et al. (2025), which turns out to be an instance of Halpern iteration (Section 4). Interestingly, in graph neural networks, APPNP (Klicpera et al., 2019) also uses input injection, where it is motivated by personalization, which our Halpern connection provides a precise characterization of by showing that the fixed point is input-dependent. The main difference from the recurrent-depth literature is that APPNP does not use truncated backpropagation.

Lipschitz neural networks Our approach of deriving weight constraints from the architecture is similar in spirit to the modular norm (Large et al., 2024; Bernstein and Newhouse, 2024), and we extend this reasoning to squared ReLU, SwiGLU, gated attention and gated SSMs. We depart from the modular norm work by not focusing on smoothness, since smoothness with respect to the ∞-norm implies smoothness with respect to every p-norm with the same Lipschitz constant. Instead, we take inspiration from the rates of Pethick (2026, Cor. 3.1), where the leading term does not involve smoothness at all but depends on the initial distance to the solution $x^\star$ in the chosen norm $\|x^0 - x^\star\|_p$, which we use to construct informative synthetic examples in Section 7.1.

We build on a long line of work on Lipschitz neural networks. For MLPs, Gouk et al. (2021) bound the Lipschitz constant by composing layerwise operator norms and enforce spectral constraints, as we do in Theorem 5.1. A close work to our fixed-point perspective is Combettes and Pesquet (2020) which analyzes neural networks as averaged, and hence nonexpansive, operators. Meunier et al. (2022) also consider nonexpansive layers by deriving them from convex potentials. Regarding attention, dot-product attention is not globally Lipschitz (Kim et al., 2021) but was shown to be Lipschitz on bounded inputs (Castin et al., 2024). Our construction specifically relies on Sinkformers (Sander et al., 2022), which use a doubly stochastic attention matrix, so that mixing the tokens is nonexpansive. We use this property to make the entire attention block Lipschitz, by additionally constraining the weights and normalizing the input. Our arguments additionally rely on the Lipschitz properties of the normalization layers, specifically LayerNorm, which with gain γ and stabilizer $\epsilon > 0$ was shown to be $|\gamma|/\sqrt{\epsilon}$-Lipschitz in Sato and Iiduka (2026); Emadi (2026).

Most existing certificates only constrain spectral norms. Our certificates differ by also requiring row-norm constraints, either alone for gates or jointly with spectral constraints for SwiGLU. We show that even simply changing the activation from ReLU to squared ReLU departs from the spectral norm as the appropriate weight constraint.

Row norms Row-wise normalization have recently been used extensively for deep learning. Loshchilov et al. (2025) normalizes the rows of the matrices reading from the residual stream in nGPT, The Scion framework (Pethick et al., 2025c) proposed a row-norm-based update, and NorMuon (Li et al., 2025), Aurora (Dewulf et al., 2026) and Muown (Lion et al., 2026) combine the spectral update of Muon with row-wise normalization or control. Lau and Su (2026) derive update directions from the symmetries of each layer, which for SwiGLU yielded a family of row-norm, one-sided spectral and hybrid updates for all three projections. In contrast, our Lipschitz perspective determines which matrices should be constrained, and in which norms, specifically, jointly in row and spectral norm for the gate and value matrices, in spectral norm for the output matrix, and in row norm alone for attention gates.

3 Problem and desiderata

We start by formulating a set of reasonable desiderata for our recurrent procedure. What we will find is that *input injection* commonly used in recurrent models such as Huginn-0125 (Geiping et al., 2025) can satisfy these requirements. Abstractly the recurrent architecture in Geiping et al. (2025) takes the following form:

$$y = \Big(\underbrace{\mathbb{D}}_{\text{decoder}} \circ \underbrace{\mathbb{G}^k}_{k \text{ recurrent steps}} \circ \underbrace{\mathbb{E}}_{\text{encoder}} \Big)(x).$$

where x is the input query, the encoder $\mathbb{E}$ (also called the Prelude) maps it to an encoded query $u = \mathbb{E}(x)$, and the decoder $\mathbb{D}$ (also called the Coda block) maps back the final latent state to the outputted answer y. The notation $\mathbb{G}^k$ denotes k successive recurrent updates.

Our work focuses on the recurrence $\mathbb{G}$ and asks how to choose this iterative process as a fixed point solver of some underlying operator. Specifically, we are given a single-valued operator parameterized by θ,

$$\mathbb{T}_\theta : \mathcal{H} \to \mathcal{H}$$

where $\mathcal{H}$ is a finite-dimensional Hilbert space and $\mathbb{T}_\theta$ has fixed-point set

$$\mathcal{C}_\theta := \mathrm{Fix}(\mathbb{T}_\theta) := \{z \in \mathcal{H} : \mathbb{T}_\theta(z) = z\}.$$

Given an encoded query represented by $u \in \mathcal{H}$, we seek a recurrent procedure

$$z^{k+1} = \mathbb{G}_{k,\theta,u}(z^k) \tag{1}$$

which should be able to return a potentially distinct solution from $\mathcal{C}_\theta$ based on each query u. Furthermore, the parameters θ should be trainable with truncated backpropagation. The following desiderata make these criteria precise.

Desideratum 1 (Existence and multiple solutions)**.** The fixed-point set $\mathcal{C}_\theta$ is nonempty, and the operator class permits $\mathcal{C}_\theta$ to contain more than one point.

If $\mathcal{C}_\theta$ were a singleton then the underlying operator $\mathbb{T}_\theta$ would not be rich enough to encode different solutions for different queries. So we need a structural assumption on the operator that permits multiple solutions.

Definition 3.1. *The operator $\mathbb{T}$ is L-Lipschitz if*

$$\|\mathbb{T}(x) - \mathbb{T}(y)\| \le L\|x - y\| \qquad \text{for all } x, y \in \mathcal{H}. \tag{2}$$

We call the operator *contractive* if (2) holds with some $L < 1$, *nonexpansive* if it holds with $L = 1$, and *expansive* if it holds only with $L > 1$. A contraction has a unique fixed point, so we will require that the operator is at least nonexpansive. See Ryu et al. (2022) for an excellent overview.

Desideratum 2 (Closest-point selection)**.** For every input $u \in \mathcal{H}$ and initialization $z^0 \in \mathcal{H}$, the iterates generated by $z^{k+1} = \mathbb{G}_{k,\theta,u}(z^k)$ satisfy

$$z^k \longrightarrow \mathbb{P}_{\mathcal{C}_\theta}(u) := \arg\min_{z \in \mathcal{C}_\theta} \tfrac{1}{2}\|z - u\|^2.$$

Through this desideratum, different queries can select different solutions from the fixed-point set of one shared operator $\mathbb{T}_\theta$ (Figure 1). The specific choice of the *closest-point* is a natural choice since it also provides robustness guarantees as we will return to in Section 4.3.

Finally, we would like the recurrent procedure to be efficient to train, in the sense that we would *not* want to backpropagate through all K steps of the recurrence. Instead, let $(\partial_\theta z^K)_{\mathrm{TBPTT}(H)}$ be obtained by differentiating only the last H updates with z^{K-H} held constant. We want that this approximation is not too far from the true gradient obtained by full backpropagation even when H only very loosely depends on K.

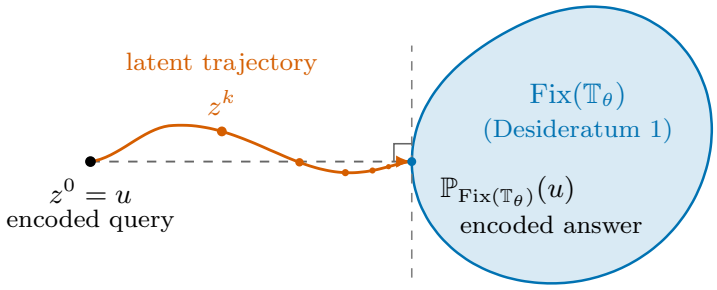

Figure 1: Closest-point selection (Desideratum 2). The fixed-point set $\mathcal{C}_\theta = \mathrm{Fix}(\mathbb{T}_\theta)$ of a nonexpansive operator is closed and convex and may contain many solutions. Started from the query/anchor u, the iterates z^k converge to the closest fixed point $\mathbb{P}_{\mathcal{C}_\theta}(u)$.

Desideratum 3 (Polylogarithmic-window TBPTT)**.** For a differentiable recurrence (1) with u and z^0 independent of θ, a backward window of size $H = \mathrm{polylog}(K, 1/\varepsilon)$ suffices to ensure

$$\left\|\partial_\theta z^K - (\partial_\theta z^K)_{\mathrm{TBPTT}(H)}\right\| \le \varepsilon.$$

Here the norm is the induced operator norm, and polylog means at most polynomial growth in $\log K$ and $\log(1/\varepsilon)$. The window is capped at K, so shorter depths use full backpropagation.

4 Halpern iteration for recurrent depth

We seek an iteration that satisfies Desiderata 1 to 3. It turns out that the classical method of Halpern iteration (Halpern, 1967) provides a way to meet these requirements and that Halpern iteration in turn is exactly what is known as *input injection*.

Halpern iteration For an anchor $u \in \mathcal{H}$ and a schedule $\beta_k \in [0, 1]$, Halpern iteration is

$$z^{k+1} = \mathbb{H}_{\beta_k, u}(z^k) := \beta_k u + (1 - \beta_k)\mathbb{T}_\theta(z^k). \qquad \text{(Halpern)}$$

The update is celebrated, since with the standard schedule $\beta_k = 1/(k+2)$ and $z^0 = u$, Halpern iteration achieves an accelerated $O(1/k)$ fixed-point residual rate for nonexpansive $\mathbb{T}_\theta$ (Sabach and Shtern, 2017; Lieder, 2021), which is worst-case optimal (Park and Ryu, 2022). In this work we are interested in Halpern for different reasons. In particular, since we seek compatibility with truncated backpropagation, we instead consider constant and logarithmic stepsize schedules.

Let us first remark on where this update comes from to get a better understanding.

Fixed-step Halpern evaluates a resolvent For a fixed schedule $\beta_k \equiv \beta \in (0, 1)$, let $\lambda := (1 - \beta)/\beta$ and define the displacement map $\mathbb{D}_\theta := \mathbb{I} - \mathbb{T}_\theta$. For $z_\beta \in \mathrm{Fix}(\mathbb{H}_{\beta,u})$, we have,

$$z_\beta = \mathbb{H}_{\beta,u}(z_\beta) \quad \Longleftrightarrow \quad z_\beta + \lambda \mathbb{D}_\theta(z_\beta) = u \quad \Longleftrightarrow \quad z_\beta = \mathbb{J}_\lambda(u),$$

with the resolvent $\mathbb{J}_\lambda := (\mathbb{I} + \lambda\mathbb{D}_\theta)^{-1}$. So the limit of Halpern with fixed $\beta_k \equiv \beta$ evaluates one application of the resolvent. This relationship already illuminates why Halpern is such a useful iteration for us. For nonexpansive $\mathbb{T}_\theta$, $\mathbb{J}_\lambda(u) \to \mathbb{P}_{\mathrm{Fix}(\mathbb{T}_\theta)}(u)$ as $\lambda \to \infty$ (Bauschke and Combettes, 2017), which is precisely Desideratum 2.

Additive input injection is fixed-step Halpern It turns out that practitioners in deep learning have converged to the same update rule. Recurrent architectures repeatedly inject the

Method	Converges to $\mathrm{Fix}(\mathbb{T}_\theta)$?	Backward horizon	Operator
Fixed-β Halpern	No†	Depth-independent	Lip $< \frac{1}{1-\beta}$
Logarithmic Halpern	Yes	Polylogarithmic in K	Nonexpansive

† Unless trivially $u \in \mathrm{Fix}(\mathbb{T}_\theta)$.

Table 1: Fixed-point selection, truncated backpropagation horizons and admissible operators of the two Halpern updates.

original input into the latent state, as in Bansal et al. (2022); Geiping et al. (2025). For additive injection, write the learned recurrent block as $\mathbb{S}_\theta$:

$$x^{k+1} = \mathbb{S}_\theta(u + x^k), \qquad x^0 \sim \mathcal{N}(0, \sigma_0^2 I).$$

Define $z^k := (u + x^k)/2$ and $\mathbb{T}_\theta(z) := \mathbb{S}_\theta(2z)$. Then

$$z^{k+1} = \tfrac{1}{2}u + \tfrac{1}{2}\mathbb{S}_\theta(2z^k) = \tfrac{1}{2}u + \tfrac{1}{2}\mathbb{T}_\theta(z^k) = \mathbb{H}_{1/2,u}(z^k).$$

Thus additive input injection is exactly Halpern with $\beta_k \equiv 1/2$ in the transformed state z^k and initialization $z^0 = (u + x^0)/2$.

To cast light on input injection we now study Halpern iteration with constant and vanishing stepsize (*c.f.* Table 1). Proofs of the results in this section can be found in Appendix A.

4.1 Fixed-step Halpern: two out of three

One recent work that also studies Halpern with a fixed stepsize is Diakonikolas (2025) which uses it to develop a restart strategy. Here we study the implications of fixed stepsize Halpern for recurrent models and truncated backpropagation.

The primary reason why Halpern with fixed stepsize is so useful is that the anchoring makes the update a contraction even when the underlying operator is mildly expansive. This in turn is what enables truncated backpropagation as the following theorem makes precise.

Proposition 4.1 (Fixed-step Halpern)**.** *Consider Halpern with $\beta_k \equiv \beta \in (0,1)$, and suppose $\mathbb{T}_\theta$ is L-Lipschitz with*

$$q := (1-\beta)L < 1.$$

Note that the operator may be expansive ($L > 1$) provided $L < 1/(1-\beta)$. Then:

(i) Contraction. *The update $\mathbb{H}_{\beta,u}$ is a contraction with modulus q, so it has a unique fixed point z_β, and $\|z^k - z_\beta\| \le q^k \|z^0 - z_\beta\|$ for every initialization.*

(ii) Depth-independent TBPTT window (Desideratum 3). *If the iteration is differentiable along its trajectory, u and z^0 are independent of θ, and $\|\partial_\theta \mathbb{H}_{\beta,u}(z^k)\| \le B$ uniformly in k, then, for $0 \le H \le K$,*

$$\left\|\partial_\theta z^K - (\partial_\theta z^K)_{\mathrm{TBPTT}(H)}\right\| \le \tfrac{B}{1-q} q^H.$$

Thus, a window $H = O(\log(1/\varepsilon))$, capped at K, suffices for error at most ε.

(iii) Approximate fixed-point selection. *The selected equilibrium z_β has fixed-point residual*

$$\|z_\beta - \mathbb{T}_\theta(z_\beta)\| = \tfrac{\beta}{1-\beta}\|u - z_\beta\|.$$

If $\mathbb{T}_\theta$ is nonexpansive with nonempty $\mathcal{C}_\theta = \mathrm{Fix}(\mathbb{T}_\theta)$, this residual is at most $\frac{\beta}{1-\beta}\,\mathrm{dist}(u, \mathcal{C}_\theta)$. It also retains the projection's variational inequality, $\langle u - z_\beta, c - z_\beta\rangle \le 0$ for every $c \in \mathcal{C}_\theta$. It is an exact fixed point if and only if $u \in \mathrm{Fix}(\mathbb{T}_\theta)$.

Tradeoff There is a clear tradeoff that becomes apparent from this result. Proposition 4.1 needs only $\mathrm{Lip}(\mathbb{T}_\theta) < 1/(1-\beta)$, e.g., $L < 2$ at $\beta = 1/2$. A larger anchor weight thus admits a more expansive operator and a shorter TBPTT window, at the price of a larger solution bias. Fixed-step Halpern can therefore satisfy Desiderata 1 and 3, but it fails Desideratum 2, since the iterates converge to a biased solution. Next we resolve this issue by decaying β_k.

4.2 Slowly vanishing Halpern

Classical Halpern iteration removes the convergence bias by allowing the anchor weight to vanish. However, the usual $\beta_k = 1/(k+2)$ is too aggressive to be compatible with TBPTT. We adopt a logarithmic schedule below, which makes the backpropagation window grow only polylogarithmically with the forward depth. Here, log denotes the natural logarithm.

Proposition 4.2 (Logarithmic-step Halpern)**.** *Assume that $\mathbb{T}_\theta$ is nonexpansive with nonempty $\mathcal{C}_\theta = \mathrm{Fix}(\mathbb{T}_\theta)$ (Desideratum 1). Then Halpern with $\beta_k := \frac{1}{\log(k+3)}$ has the following properties.*

(i) Closest-point selection (Desideratum 2). *For every $z^0 \in \mathcal{H}$, $z^k \longrightarrow \mathbb{P}_{\mathcal{C}_\theta}(u)$.*

(ii) Polylogarithmic-window TBPTT (Desideratum 3). *Suppose the recurrence is differentiable along its trajectory, u and z^0 are independent of θ, and $\sup_k \|\partial_\theta \mathbb{T}_\theta(z^k)\| \le M$. Then, for $0 \le H < K$,*

$$\left\|\partial_\theta z^K - (\partial_\theta z^K)_{\mathrm{TBPTT}(H)}\right\| \le M \log(K+2) \exp\left(-\tfrac{H}{\log(K+2)}\right). \tag{3}$$

Thus, a window $H = O(\log K\,[\log\log K + \log(1/\varepsilon)])$ suffices for absolute error at most ε.

Tradeoff Thus, slowly vanishing Halpern can satisfy all of Desiderata 1 to 3, with the price being that the operator $\mathbb{T}_\theta$ is restricted to be nonexpansive.

4.3 Query robustness

As required by Desideratum 2, slowly vanishing Halpern selects the closest fixed point $\mathbb{P}_{\mathcal{C}_\theta}(u)$. Since projection onto a closed convex set is (firmly) nonexpansive, $\|\mathbb{P}_{\mathcal{C}_\theta}(u) - \mathbb{P}_{\mathcal{C}_\theta}(\widetilde{u})\| \le \|u - \widetilde{u}\|$. This provides a robustness guarantee stating that perturbing the query u will only perturb the solution by as much. We can say something stronger for the specific case of Halpern. The following theorem extends this robustness to every finite iterate and every schedule β_k.

Theorem 4.3 (Query robustness of Halpern)**.** *Let $\mathbb{T}_\theta$ be L-Lipschitz, $\beta_k \in [0,1]$, $\beta_{\min} := \inf_k \beta_k > 1 - \frac{1}{L}$, and*

$$C := 1 \quad \text{if } L \le 1, \qquad C := \tfrac{\beta_{\min}}{1-(1-\beta_{\min})L} \quad \text{if } L > 1. \tag{4}$$

If (z^k) and $(\widetilde{z}^k)$ follow $z^{k+1} = \mathbb{H}_{\beta_k,u}(z^k)$ and $\widetilde{z}^{k+1} = \mathbb{H}_{\beta_k,\widetilde{u}}(\widetilde{z}^k)$, and $\|u - \widetilde{u}\|, \|z^0 - \widetilde{z}^0\| \le \varepsilon$, then

$$\|z^k - \widetilde{z}^k\| \le C\varepsilon \qquad \text{for every } k \ge 0. \tag{5}$$

The theorem applies even to expansive operators ($L > 1$) for a possibly vanishing schedule as long as there is a floor $\beta_k \ge \beta_{\min} > 1 - 1/L$. This floor is necessary since the constant C is sharp for $\beta_k \equiv \beta$, as realized by the simple example $\mathbb{T}_\theta = L\mathbb{I}$.

5 A Lipschitz neural parameterization

We have satisfied the stated desiderata, but we have assumed away one major obstacle, namely ensuring that the Lipschitz constant of the operator $\mathbb{T}_\theta$ is controlled. Specifically, Proposition 4.1 requires $L < 1/(1-\beta)$. This section is devoted to constructing such Lipschitz certificates along with ensuring existence of fixed points for common architectural components such as MLPs, SwiGLU, and attention. We want to highlight that these results are not only relevant for recurrent models, but will inform us on how to train general deep models even without weight-tying. Proofs and supplementary results for this section are collected in Appendix B.

Normalization layers Normalization layers can be made nonexpansive with appropriate parameterization. For $x \in \mathbb{R}^d$, write $r := \mathrm{rms}(x) = \|x\|/\sqrt{d}$. For $0 < a < \tau$, we use

$$\mathbb{N}_{\mathrm{RMS}}(x) := \frac{\tau x}{\sqrt{\tau^2 + r^2}}, \qquad \mathbb{N}_{\tanh}(x) := \frac{a + (\tau - a)\tanh\left(\frac{(r-a)_+}{\tau - a}\right)}{\max\{r, a\}}\, x, \tag{6}$$

applied to each row of a sequence $X \in \mathbb{R}^{n \times d}$. $\mathbb{N}_{\mathrm{RMS}}$ is simply RMSNorm (Zhang and Sennrich, 2019) with scalar gain $g = \tau$ and stabilizer $\epsilon = \tau^2$. For general $g > 0$ and $\epsilon > 0$, RMSNorm has Lipschitz constant $g/\sqrt{\epsilon}$ (Chen et al., 2026, Lemma A.3), so it is nonexpansive when $g/\sqrt{\epsilon} \le 1$. We provide as an alternative, $\mathbb{N}_{\tanh}$ (which is the identity up to RMS a and smoothly saturates at RMS τ) because its identity region permits multiple fixed points when composed with other nonexpansive operators, whereas compositions involving $\mathbb{N}_{\mathrm{RMS}}$ have at most one. Both $\mathbb{N}_{\mathrm{RMS}}$ and $\mathbb{N}_{\tanh}$ are instances of the more general *normalization layer* $\mathbb{N}$ defined in Definition B.1 in Appendix B with parameter $\omega = \tau/2$ for $\mathbb{N}_{\mathrm{RMS}}$ and $a \le \omega \le \tau$ for $\mathbb{N}_{\tanh}$.

MLP As a warmup we start with a simple MLP using an L_σ-Lipschitz activation function, such as ReLU for which $L_\sigma = 1$.

Theorem 5.1 (Lipschitz MLP block)**.** *Let $W_1 : \mathbb{R}^d \to \mathbb{R}^m$ and $W_2 : \mathbb{R}^m \to \mathbb{R}^d$ be linear maps, let $b_1 \in \mathbb{R}^m$ and $b_2 \in \mathbb{R}^d$, and suppose that $\sigma : \mathbb{R}^m \to \mathbb{R}^m$ is L_σ-Lipschitz in the Euclidean norm. Define the raw MLP*

$$\mathbb{F}_{\mathrm{MLP}}(x) := W_2\sigma(W_1 x + b_1) + b_2 \tag{7}$$

and, for a normalization layer $\mathbb{N}$, its input- and output-normalized variants

$$\mathbb{M}_{\mathrm{in}} := \mathbb{F}_{\mathrm{MLP}} \circ \mathbb{N}, \qquad \mathbb{M}_{\mathrm{out}} := \mathbb{N} \circ \mathbb{F}_{\mathrm{MLP}}. \tag{8}$$

Then

$$\mathrm{Lip}(\mathbb{F}_{\mathrm{MLP}}),\ \mathrm{Lip}(\mathbb{M}_{\mathrm{in}}),\ \mathrm{Lip}(\mathbb{M}_{\mathrm{out}}) \le \|W_2\|_{\mathrm{op}} L_\sigma \|W_1\|_{\mathrm{op}}. \tag{9}$$

Remark 5.2 (Placement of the MLP normalization). The bound (9) already holds for the raw MLP, without normalization. However, the bounded range enforced by normalization layers is used for showing existence of fixed points (Appendix B.5).

SwiGLU and ReLU^2 An activation function that leads to much more interesting conclusions is the SwiGLU block (Shazeer, 2020), which has been widely adopted in the field. Due to the gating mechanism there is an elementwise quadratic interaction of the input which in turns leads to a Jacobian that grows linearly in the input norm. This will have two important consequences, one being that input normalization (also known as pre-norm) becomes essential, and second, that it naturally leads to a joint row-norm and spectral constraint on specific matrices. Squared ReLU shares this structure and thus leads to a similar bound involving both row and spectral norms, which additionally enjoys a particularly simple argument (see Theorem B.5).

Block	Normalize	Constrained matrices		Lipschitz bound
		Spectral ($2 \to 2$)	Row ($2 \to \infty$)	
MLP (Theorem 5.1)	None†	W_1, W_2	—	$L_\sigma s_1 s_2$
ReLU^2 MLP (Theorem B.5)	Input	W_1, W_2	W_1	$2\sqrt{d}\,\omega s_1 s_2 r_1$
SwiGLU (Theorem 5.3)	Input	W_o, W_g, W_v	W_g, W_v	$\sqrt{d}\,\omega s_o(L_{\mathrm{SiLU}} r_v s_g + r_g s_v)$
Attention (Theorem B.7)	Input	W_Q, W_K, W_V, W_O	—	$(1 + 2Be^{4B}) s_V s_O$
Gated attention (Theorem B.9)	Input	W_O	W_a	$c s_O (L_U + M r_a/4)$
Gated SSM (Theorem B.10)	Input	W_V, W_O	W_a	$\frac{s_O s_V}{1-\rho}(1 + \frac{\sqrt{d}\,\tau r_a}{4(1-\rho)})$

† The MLP needs no normalization for its Lipschitz bound, but the boundedness-based existence guarantee relies on input or output normalization (Theorem B.11).

Table 2: Lipschitz bounds of the blocks. Let $s_i := \|W_i\|_{\mathrm{op}}$ and $r_i := \|W_i\|_{2\to\infty}$.

Theorem 5.3 (Lipschitz SwiGLU block)**.** *Let $W_g, W_v : \mathbb{R}^d \to \mathbb{R}^m$ and $W_o : \mathbb{R}^m \to \mathbb{R}^d$ be linear maps, and let $\mathbb{N}$ be a normalization layer. Write $\mathrm{SiLU}(t) = t/(1+e^{-t})$, applied coordinatewise, and let $L_{\mathrm{SiLU}} := \sup_t |\mathrm{SiLU}'(t)| \approx 1.1$. Define the maximum row-norm $\|W\|_{2\to\infty} := \max_i \|W_{i,:}\|_2$ and the input-normalized SwiGLU map*

$$\mathbb{M}_{\mathrm{SwiGLU}}(x) := W_o[\mathrm{SiLU}(W_g\mathbb{N}(x)) \odot W_v\mathbb{N}(x)]\,. \tag{10}$$

Then

$$\mathrm{Lip}(\mathbb{M}_{\mathrm{SwiGLU}}) \le \sqrt{d}\,\omega\|W_o\|_{\mathrm{op}}\left[L_{\mathrm{SiLU}}\|W_v\|_{2\to\infty}\|W_g\|_{\mathrm{op}} + \|W_g\|_{2\to\infty}\|W_v\|_{\mathrm{op}}\right]. \tag{11}$$

Remark 5.4 (Spectral and rownorm split for SwiGLU). There is freedom in how we allocate the Lipschitz budget. If $\|W_o\|_{\mathrm{op}} \le 1$ and W_g, W_v have spectral norms at most s and row norms at most r, then $\mathrm{Lip}(\mathbb{M}_{\mathrm{SwiGLU}}) \le (1 + L_{\mathrm{SiLU}})\sqrt{d}\,\omega\, rs$. The certificate depends on the radii only through the budget $B := rs$, which nonexpansiveness fixes at $B \le 1/((1+L_{\mathrm{SiLU}})\sqrt{d}\,\omega)$, with the ratio $\kappa := r/s \in (0, 1]$ being free. To choose a split, consider what we will call a *saturated* matrix $W \in \mathbb{R}^{m\times d}$ with $m \ge d$, where all d singular values equal to s, and all m row norms equal to r. Then $mr^2 = \|W\|_F^2 = ds^2$, giving $r/s = \sqrt{d/m}$. For a wide matrix, $m < d$, all m singular values equal s, so $WW^\top = s^2 I_m$ and the rows are orthogonal with norm exactly $r = s$. We therefore recommend $\kappa = \min\{1, \sqrt{d/m}\}$, as in Algorithm 1, which makes the row and spectral bounds simultaneously tight for this class of fully saturated matrices (see Section 7.1).

Insight 5.1. SwiGLU blocks and MLP blocks using ReLU^2 should optimize over the intersection of the spectral and row norm unit balls for which we propose the Sparrow optimizer (*c.f.* Section 6).

Attention We use Sinkhorn-based attention (Sander et al., 2022), whose Lipschitz certificate is given in Theorem B.7 (Appendix B.3) revealing spectral matrix constraints. Similarly to the SwiGLU block, input normalization turns out to be crucial for attention. However, note that no QK-normalization is needed for the Lipschitz certificate. The result can be extended to softmax attention by adapting Castin et al. (2024, Theorem 3.5), where the importance of pre-norm is already apparent since the result relies on boundedness of the input.

Gated architectures In the appendix we additionally derive Lipschitz certificates for gated attention and gated SSMs in Appendices B.4.1 and B.4.2, which naturally leads to a constraint on the row norms of the gate weights without requiring a spectral bound. Interestingly, in a recent

hero run, Marin reported attention-gate norm growth (Marin Community, 2026b) and intervened by adding decoupled weight decay to the attention-gate (Marin Community, 2026a), while using the spectral-based optimizer MuonH. Our theory suggests instead optimizing the gate matrices with a row-norm-based optimizer, such as Scion (RowNorm) (Pethick et al., 2025c). This is the attention-gate branch of Sparrow, which Section 7.2 shows helps when the gate matrix is tall.

Combining attention and MLP blocks Transformer architectures evaluate attention and MLP blocks either serially or in parallel. Suppressing normalization and parameters, the serial block with residual connections as in ResNet (He et al., 2016) and the parallel block used by e.g., PaLM and PAF (Chowdhery et al., 2022; Sonkar and Baraniuk, 2023) have the forms

$$\begin{aligned} \text{serial residual:} \quad & X^+ = (\mathbb{I} + \mathbb{M}) \circ (\mathbb{I} + \mathbb{A})(X), \\ \text{parallel residual:} \quad & X^+ = X + \mathbb{A}(X) + \mathbb{M}(X). \end{aligned}$$

Because of the unscaled identity residual, their Lipschitz constants are up to $(1 + L_M)(1 + L_A)$ and $1 + L_A + L_M$ when $\mathbb{A}$ and $\mathbb{M}$ are L_A- and L_M-Lipschitz, so neither is nonexpansive even when the branches are. We instead replace the additions by convex combinations.

Proposition 5.5 (Serial and parallel Lipschitz blocks)**.** *Let $\mathbb{A}_\theta, \mathbb{M}_\theta : \mathcal{H} \to \mathcal{H}$ be L_A- and L_M-Lipschitz, and let $\alpha_A, \alpha_M, \alpha \in [0, 1]$. Define*

$$\mathbb{F}_A := (1 - \alpha_A)\mathbb{I} + \alpha_A \mathbb{A}_\theta, \qquad \mathbb{F}_M := (1 - \alpha_M)\mathbb{I} + \alpha_M \mathbb{M}_\theta, \tag{12}$$

$$\mathbb{T}_{\text{ser}} := \mathbb{F}_M \circ \mathbb{F}_A, \qquad \mathbb{T}_{\text{par}} := (1 - \alpha)\mathbb{A}_\theta + \alpha \mathbb{M}_\theta. \tag{13}$$

Then $\operatorname{Lip}(\mathbb{T}_{\text{ser}}) \le (1 - \alpha_M + \alpha_M L_M)(1 - \alpha_A + \alpha_A L_A)$, *and* $\operatorname{Lip}(\mathbb{T}_{\text{par}}) \le (1 - \alpha)L_A + \alpha L_M$.

Table 2 summarizes our Lipschitz bounds for the blocks and their implied weight constraints.

Fixed-point existence Normalization layers are sufficient to ensure that a fixed point exists, since every block of the form $\mathbb{G} \circ \mathbb{N} \circ \mathbb{F}$ with continuous $\mathbb{F}, \mathbb{G}$, which covers the (normalized) MLP, SwiGLU, ReLU^2, attention and gated blocks above, has a fixed point regardless of its Lipschitz constant (Theorem B.11 in Appendix B.5).

6 Optimizing the constrained weights

Our Lipschitz certificates show that our weight matrices should be constrained in particular ways as summarized in Table 2. This section deals with how our optimizer takes these constraints into account. In particular, we will focus on computing the linear maximization oracle (LMO) and projection:

$$\mathsf{lmo}_{\mathcal{K}}(G) := \arg\max_{Z \in \mathcal{K}} \langle G, Z \rangle_F, \qquad \mathbb{P}_{\mathcal{K}}(Y) := \arg\min_{Z \in \mathcal{K}} \tfrac{1}{2}\|Z - Y\|_F^2. \tag{14}$$

Both minimize $q(W) := \frac{a}{2}\|W\|_F^2 - \langle H, W \rangle_F$ over $\mathcal{K}$, with $(a, H) = (0, G)$ for the LMO and $(a, H) = (1, Y)$ for the projection.

Single-constraint oracles To set up notation, consider the spectral and row-norm oracles under the single constraints $\|W\|_{\text{op}} \le s$ and $\|W\|_{2\to\infty} \le r$, respectively. For a compact SVD $X = U \operatorname{diag}(\sigma_j) V^\top$,

$$\begin{aligned} \text{LMO:} \quad & \mathsf{Spectral}_s(X) = sUV^\top, && \mathsf{Row}_r(X)_{i,:} = r\frac{X_{i,:}}{\|X_{i,:}\|_2}, \\ \text{projection:} \quad & \mathsf{SpectralClip}_s(X) = U \operatorname{diag}(\min\{\sigma_j, s\}) V^\top, && \mathsf{RowClip}_r(X)_{i,:} = \frac{X_{i,:}}{\max\{1, \|X_{i,:}\|_2 / r\}}, \end{aligned}$$

Algorithm 1 Sparrow update direction. It applies $\mathsf{Spectral}_s$ at most once and is always feasible.

Require: gradient estimate $G \in \mathbb{R}^{m\times d}$, radii s and $r = s\min\{1, \sqrt{d/m}\}$ (see Remark 5.4)
Return: depending on the (hidden) parameter group of G:

1: SwiGLU gate and value matrices $(m > d)^\dagger$:	$(\mathsf{RowClip}_r \circ \mathsf{Spectral}_s \circ \mathsf{Row}_r)(G)$
2: Attention-gate/SSM-gate matrices:	$\mathsf{Row}_r(G)$
3: All other matrices:	$\mathsf{Spectral}_s(G)$

†Recommended solver, see Section 6.1 for alternative solvers and Algorithm 2 for the general case.

6.1 Joint row-norm and spectral-norm constraints

The constraint that requires extra care is the SwiGLU bound (11) constraining W_g, W_v to

$$\mathcal{K}_{r,s} := \left\{W \in \mathbb{R}^{m\times d} : \|W\|_{2\to\infty} \le r,\ \|W\|_{\mathrm{op}} \le s\right\}, \qquad r, s > 0. \tag{15}$$

Since $\|W\|_{2\to\infty} \le \|W\|_{\mathrm{op}} \le \sqrt{m}\,\|W\|_{2\to\infty}$, both constraints are active only when $r < s < \sqrt{m}\,r$. For $m \ge d$, the recommended split $r/s = \sqrt{d/m}$ in Remark 5.4 makes both constraints nonredundant when $m > d > 1$. At the endpoints $m = d$ and $d = 1$, the spectral bound implies the row bound and the row bound implies the spectral bound, respectively. In other words, for square matrices the row bound is redundant.

Sparrow Since every method needs a name, we will call an (approximate) solution to the LMO and projection oracles Sparrow which is short for *Spectral and row-norm restricted optimization of weights*. Algorithm 1 gives the LMO update we recommend based on the weight constraints summarized in Table 2 and the approximation of the joint constraint developed below. In practice, the LMO can be combined with standard techniques such as a Nesterov momentum based gradient estimator, weight decay and weight normalization (Appendix C.1) as done in our accompanying PyTorch and Jax implementation. Sparrow provides a drop-in replacement of Muon/MuonH that leads to improvements without any retuning.

Coarse feasible approximation Neither oracle of $\mathcal{K}_{r,s}$ has a closed form. A cheap candidate applies the oracle of one (or two) constraints and then clips with respect to the other constraint:

$$\begin{aligned}
\text{LMO:}\quad & U_1 = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s)(G), \quad U_2 = (\mathsf{SpectralClip}_s \circ \mathsf{Row}_r)(G), \\
& U_3 = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s \circ \mathsf{Row}_r)(G), \\
\text{projection:}\quad & V_1 = (\mathsf{RowClip}_r \circ \mathsf{SpectralClip}_s)(Y), \quad V_2 = (\mathsf{SpectralClip}_s \circ \mathsf{RowClip}_r)(Y).
\end{aligned} \tag{16}$$

Alternatively, one can keep the candidate with the smallest value of q, which costs $O(md)$ to check. Importantly, the output is feasible, since row clipping scales rows by factors at most one, so it cannot increase the spectral norm, and spectral clipping cannot increase any row norm. In contrast, Aurora is only feasible for the spectral norm and not necessarily for the row norm.

As we shall see next, these coarse approximations can be seen as performing one step of an iterate splitting algorithm for solving the dual problem of (14) with $\mathcal{K} = \mathcal{K}_{r,s}$.

Dual problem and forward-backward splitting Since $\mathcal{K}_{r,s}$ is the intersection of the balls $\{\|Z\|_{2\to\infty} \le r\}$ and $\{\|Z\|_{\mathrm{op}} \le s\}$, whose support functions are $r\|\cdot\|_{2,1}$ and $s\|\cdot\|_{\mathrm{nuc}}$ with $\|A\|_{2,1} := \sum_i \|A_{i,:}\|_2$, the LMO has the dual

$$\max_{Z\in\mathcal{K}_{r,s}} \langle G, Z\rangle = \min_{A+B=G} r\|A\|_{2,1} + s\|B\|_{\mathrm{nuc}}. \tag{17}$$

Forward-backward splitting on (17), reduces to taking a gradient step on $s\|G - A\|_{\mathrm{nuc}}$ followed by a proximal step on $r\|A\|_{2,1}$, which reads

$$P := \mathsf{Spectral}_s(G - A), \qquad A \leftarrow \mathsf{RowShrink}_{\gamma r}(A + \gamma P), \qquad Z := \mathsf{RowClip}_r(P), \tag{FB}$$

where $\mathsf{RowShrink}_t := I - \mathsf{RowClip}_t$ is the proximal map of $t\|\cdot\|_{2,1}$ (Moreau's decomposition). An intelligent initialization is $A = \mathsf{RowShrink}_c(G)$, since this is the minimizer of (17) with $s\|B\|_{\text{nuc}}$ replaced by the quadratic $\frac{r}{2c}\|B\|_F^2$. Started from $A = \mathsf{RowShrink}_c(G)$, so that $G - A = \mathsf{RowClip}_c(G)$, the first output is

$$Z = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s \circ \mathsf{RowClip}_c)(G).$$

This is exactly our coarse approximations U_1 for $c \to \infty$ (indeed for $c \geq \max_i \|G_{i,:}\|_2$) and U_3 for $c \to 0$ (indeed for $c \leq \min_i \|G_{i,:}\|_2$, where every row of $\mathsf{RowClip}_c(G)$ has norm c and $\mathsf{Spectral}_s$ is invariant to positive scaling). See Appendix C.2 for a numerically stable implementation.

The choice of the splitting is in some sense arbitrary. We could just as well exchange the roles of the two terms in (17) by taking a gradient step on $r\|G - B\|_{2,1}$ and a proximal step on $s\|B\|_{\text{nuc}}$ with output $\mathsf{SpectralClip}_s(\mathsf{Row}_r(G - B))$, which gives U_2 with $B = 0$. The main disadvantage is computational since $\mathsf{SpectralClip}_s$ can be more expensive as we touch on later.

Iterative refinement through operator splitting If the coarse one step approximations are not accurate enough, one could iterate FB, but since $\mathsf{Spectral}_s$ is not Lipschitz, FB converges only locally and needs a small step size. Douglas-Rachford splitting (DRS), or equivalently ADMM on the dual (17), solves the joint problem for every step size, at the cost of one $\mathsf{SpectralClip}_s$ per step. It splits the objective into $f := q + \iota_{\{\|W\|_{2\to\infty} \leq r\}}$ and $g := \iota_{\{\|W\|_{\text{op}} \leq s\}}$, where ι is the convex indicator, with proximal maps

$$\text{prox}_{\gamma f}(Z) = \mathsf{RowClip}_r\left(\tfrac{Z+\gamma H}{1+\gamma a}\right), \qquad \text{prox}_{\gamma g}(Z) = \mathsf{SpectralClip}_s(Z).$$

Since the spectral step is typically inexact while the row step is exact, one can use inexact DRS (iDRS) (Pethick et al., 2025a), which allows for a relative error condition on the inexactness and provides a stopping criterion for the inner solver. For the projection, alternating projection (POCS) of the two clipping maps is another option, for which one step yields V_1, V_2 in (16).

We compare these approaches with inexact, low-rank spectral steps in Appendix C.3. In practice, however, Newton-Schulz iterations (Higham, 2008) and PolarExpress (Amsel et al., 2025) are typically used to approximate $\mathsf{Spectral}_s$ using only matrix products. $\mathsf{SpectralClip}_s$ can similarly be approximated through two approximations of $\mathsf{Spectral}_s$ (*c.f.* Appendix C.4).

7 Experiments

7.1 Robustness-accuracy pareto frontier

We create a synthetic example where we can clearly investigate when (and when not) our derived constraints can be beneficial over a purely spectrally constrained based approach like Muon. Specifically, inspired by the rates of Pethick (2026, Cor. 3.1), which depend on the distance to solution $\|x^0 - x^\star\|$ in some chosen geometry, we construct a teacher that is close in the joint row and spectral geometry, while simultaneously being far (or close) in other geometries.

Setup We train the SwiGLU block $\mathbb{M}(x) = \mathbb{G}(\mathbb{N}_{\text{RMS}}(x))$ with $\mathbb{G}(z) = W_o[\text{SiLU}(W_g z) \odot W_v z]$, $\|W_o\|_{\text{op}} \leq 1$ and $W_g, W_v \in \mathbb{R}^{m\times d}$. Since the certificate depends on W_g, W_v only through the budget $B = rs$, the product of their maximal row norm r and spectral norm s (Remark 5.4), we compare the spectral ball $\{\|W\|_{\text{op}} \leq \sqrt{B}\}$ with the joint set $\mathcal{K}_{\sqrt{B\kappa},\sqrt{B/\kappa}}$ at the same budget B, where $\kappa = r/s$ is the split.

The task is ten-class classification with labels given by the argmax of a teacher with $d = 64$ and $m = 4d$, which is roughly nonexpansive (Appendix D.1). We describe a teacher by the ratio $\kappa^\star$ of its maximal row norm to its spectral norm, and consider a Gaussian teacher whose mass is

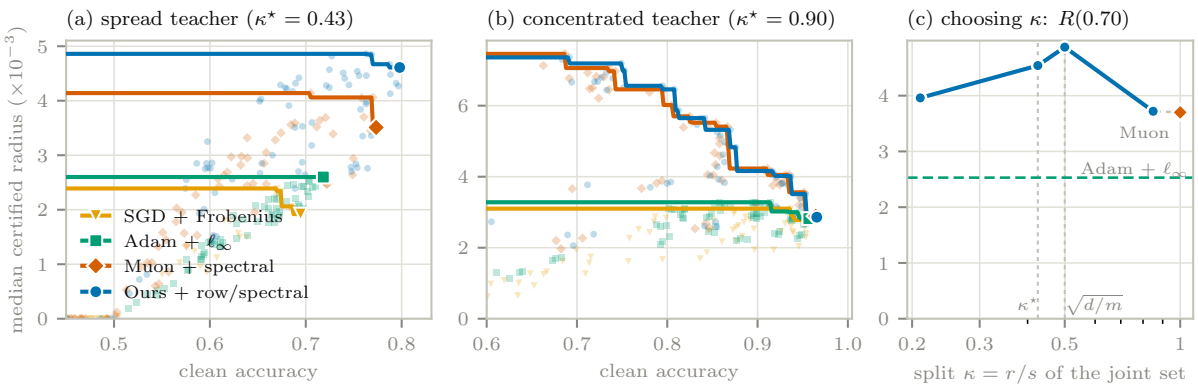

Figure 2: Robustness. Row/spectral constraints certify larger radii for a given accuracy. (a, b) Each point is one run and each line is the frontier $R_M(a)$ (3 seeds). (c) We can set the budget split as $\kappa = \sqrt{d/m}$ without needing to know the ground truth κ^*.

spread across the rows ($\kappa^\star = 0.43$) and a *concentrated* teacher with only 8 nonzero rows in W_g and W_v ($\kappa^\star = 0.90$).

We compare SGD with a Frobenius ball, Adam with an entrywise box, Muon with a spectral ball, and our update with the joint set constraint (using the approximations U_1 and V_1 from (16) for the LMO and projection, respectively), all with the same budget B. Robustness is measured by the certified radius of each test point, computed from the Lipschitz certificate of the trained weights. For each method M we sweep over step sizes and budgets for three seeds and report the Pareto frontier $R_M(a)$ defined as the largest median radius among runs with clean accuracy at least a (see Appendix D for details).

Spectral constraint We observe in Figure 2 that the spectral constraint certifies much larger radii than the Frobenius and entrywise constraints. At 70% accuracy on the spread teacher, the radius certified by Muon is 1.6× that of Adam, and SGD and Adam require 64 to 4096 times the budget of the teacher to fit the task. The budgets of the four constraints are not directly comparable, since their worst case bound can be far from a typical case. However, the comparison remains fair, since compute every radius from the trained weights instead.

The row norm constraint On the spread teacher, the joint set dominates Muon at every accuracy level, with 1.2× the radius at 70% accuracy and a higher maximal accuracy (0.793 vs. 0.771). The same holds with 3× longer training (cf. Table 5). The row constraint helps because the spectral ball cannot make use of the width as we show in Appendix D.1 by making W_g and W_v increasingly tall. On the concentrated teacher, for which $\kappa^\star$ is close to one, the two methods have a matching pareto front. In other words, the joint constraint is only useful if the problem truly requires all the parameters.

Choosing the radius ratio Since $\kappa^\star$ is not known in practice, we also vary the split κ of the joint set on the spread teacher. Figure 2c shows that $R(0.70)$ peaks near the teacher-free split $\kappa = \sqrt{d/m}$ of Remark 5.4, which is used in Sparrow. We observe that on the concentrated teacher this choice also does not hurt.

7.2 Sparrow as a drop-in replacement for Muon

Without the projection, Sparrow can be used as an update direction for unconstrained weights in place of Muon. We train the character-level Shakespeare model of nanoGPT (4 layers,

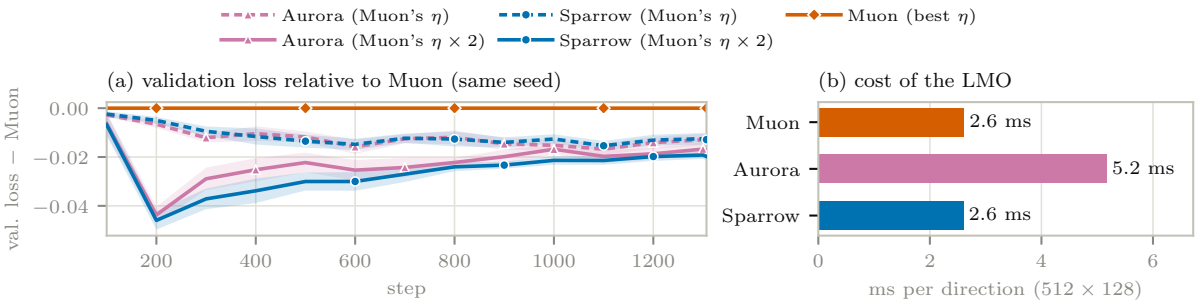

Figure 3: Sparrow lowers the validation loss at no additional cost per step. (a) Validation loss minus Muon's at the same step and seed (mean over 10 seeds with one standard error), with Muon at its best step size η and Aurora and Sparrow at $1\times$ and $2\times$ the step size of Muon. (b) The LMO in Sparrow has no wall-clock time overhead over Muon in contrast with Aurora.

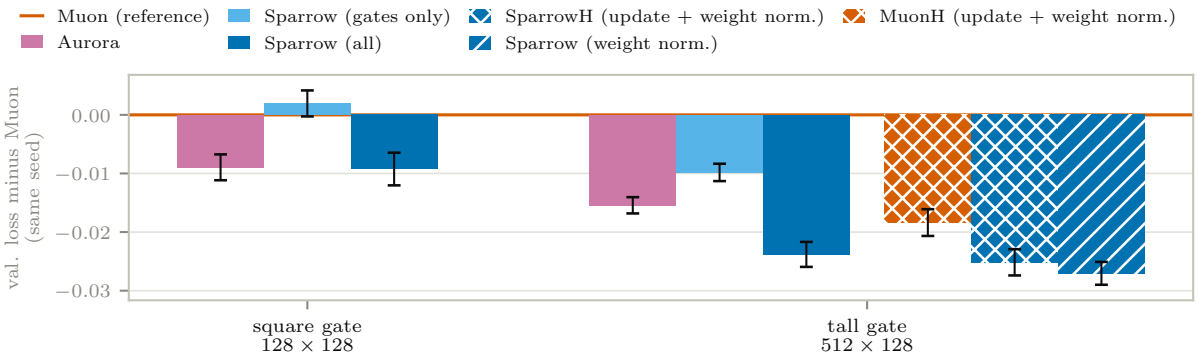

Figure 4: Gated attention: final validation loss minus that of Muon on the same seed (mean and standard error over ten seeds). The row-norm update of the gate in Sparrow helps only when the gate is tall. When combined with Frobenius normalization as in MuonH, one can ignore normalizing the update direction in Sparrow without sacrificing performance.

$d = 128$, RMSNorm, SwiGLU width $m = 4d$) for one epoch. The hidden matrices are updated as $W \leftarrow W - \eta \mathsf{lmo}(M)$, where M is the momentum, and the embeddings and head are trained with AdamW. Muon uses $\mathsf{lmo} = \mathsf{Spectral}_1$ with its best step size $\eta = 0.01$. Aurora (Dewulf et al., 2026) and Sparrow only change the lmo for W_g and W_v and otherwise use the same configuration as Muon (see Appendix D.2).

Speedup Figure 3a shows the validation loss of each method minus that of Muon on the same seed. At Muon's step size, Sparrow lowers the final validation loss on all 10 seeds, by 0.013 ± 0.002 on average, and reaches the final loss of Muon after 1220 instead of 1302 steps ($1.07\times$). At twice that step size, it lowers the loss by 0.020 ± 0.002 ($1.11\times$), on par with Aurora at the same step size (0.017 ± 0.003, $1.09\times$). A direction of Sparrow takes the same 2.6 ms as Muon's, since the additional row normalization and row clipping are essentially free in comparison with the spectral solver, whereas Aurora's two polar factors double the cost (Figure 3b). Additionally, in contrast with Aurora, the update direction of Sparrow lies in $\mathcal{K}_{r,s}$ (Appendix D.2 compares both methods with the exact LMO).

Gated attention The Lipschitz constant of the gate depends only on $\|W_a\|_{2\to\infty}$ (Theorem B.9), so Sparrow updates W_a with Row_r and drops the spectral constraint on the gate altogether. This

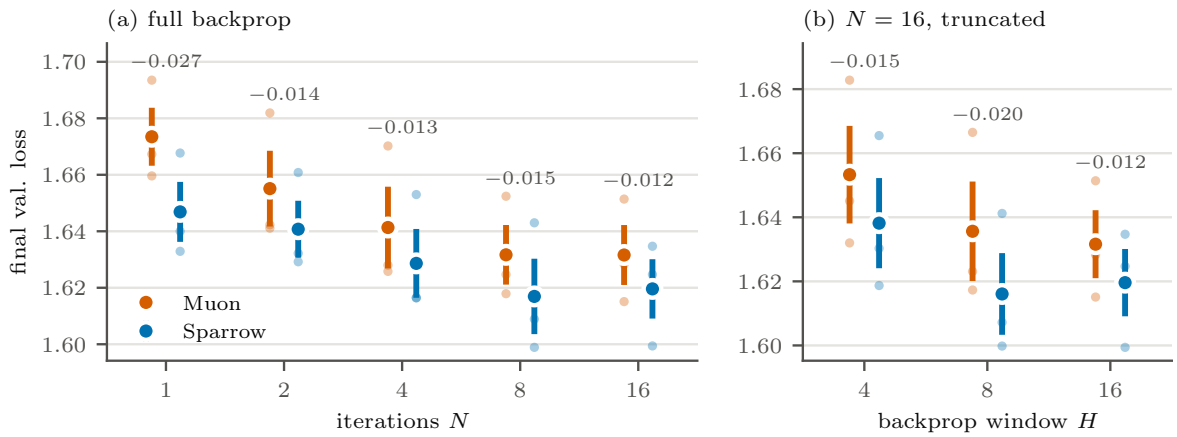

Figure 5: Looped model. Final validation loss at each method's best step size (mean ± one standard error over three seeds, the numbers indicate Sparrow minus Muon on the same seed). Sparrow consistently outperform Muon without retuning.

only makes a difference when W_a is tall, since otherwise every row of $\mathsf{Spectral}_s(G)$ already has norm s. Tall gates are indeed used in practice, e.g., Qwen3-Next gates 4096 channels at width 2048 (Qwen Team, 2025; Qiu et al., 2025). We add a gate to the model above with 4 heads of width 32, which gives a square 128×128 gate, or 128, which gives a tall 512×128 gate. We compare Muon, Aurora on every tall matrix as recommended, Sparrow on the gates only, and Sparrow on all matrices, each over ten seeds.

- *Square gate.* Using Sparrow on the gate alone has no effect ($+0.002 \pm 0.002$), as expected (Figure 4), so Sparrow is safe to use even when the gate matrices are square. The ablation also tells us that the large improvement from Sparrow (all) comes exclusively from applying Sparrow to the attention gates.
- *Tall gate.* Applying Aurora improves the most over Muon in the tall gate setting (0.015 ± 0.001). However, dropping the spectral constraint entirely as in Sparrow leads to a more than $1.5\times$ improvement compared with Aurora (0.024 ± 0.002), suggesting that the row-norm is the more appropriate constraint for the attention gates, and that Aurora should *not* be applied to these matrices despite them being tall.
- *Weight normalization.* MuonH (Wen et al., 2026) keeps each hidden matrix at its initial Frobenius norm and additionally normalizes the update direction, which lowers the loss with the tall gate by 0.018 ± 0.002.[1] Combining Sparrow with weight normalization the loss is further improved by 0.007–0.009 (± 0.002) even without using update direction normalization, which suggests that the two are complementary.

Looped models We extend the setup with the tall gate to a looped model with a prelude layer, a core layer iterated N times with fixed-step Halpern ($\beta = \frac{1}{2}$, Proposition 4.1), and a coda layer (Appendix D.3). Sparrow lowers the final validation loss for every N and every truncation window H on all three seeds, by 0.012–0.027 (Figure 5).

[1]The improvement of MuonH might be partially caused by changing away from zero initialization, which is otherwise incompatible with MuonH.

8 Perspective

Taking the Lipschitz certificate approach in order to derive efficient and robust optimization algorithms seems like an unreasonably predictive approach. Can Lipschitz certificates similarly guide optimizer design for other modules like Mixture-of-Experts (MoE)? How should we be scaling models, with optimizers that properly utilize parameters of tall matrices? Are there any model design-choices we discarded in the past because we did not optimize them appropriately?

Regarding recurrent-depth, the grounding of input injection in Halpern iteration makes it possible to ask a range of questions: Is there a recurrence that converges to a fixed point of $\mathbb{T}_\theta$, like slowly vanishing Halpern, which also tolerates operators that are not strictly nonexpansive, like fixed-step Halpern? Is it worthwhile letting go of truncated backpropagation to support a parameterization like $\beta_k = 1/(k+2)$ which enjoys faster convergence rates and thus shorter depth?

AI use statement

The desiderata, the connection between input injection and Halpern iteration, its contractive property and the consequences for truncated backpropagation, the observations surrounding robustness, and the central Lipschitz argument for squared ReLU and gating were developed by the author. Generative AI were used to extend this argument to SwiGLU and state-space models, and an LLM proposed the normalization $\mathbb{N}_{\tanh}$. The connection to forward-backward splitting and iDRS, and the practical Sparrow optimizer, were developed by the author, while the low-rank case in Appendix C.3 was proposed by generative AI. Generative AI were used extensively for formalizing arguments, writing code and producing plots. The author takes full responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Acknowledgments

The author thanks the Swiss chômage for supporting the work.

References

Ahmet Alacaoglu, Donghwan Kim, and Stephen J. Wright. Revisiting inexact fixed-point iterations for min-max problems: Stochasticity and structured nonconvexity. *arXiv preprint arXiv:2402.05071*, 2024. URL [url]. (cited on p. 2)

Noah Amsel, David Persson, Christopher Musco, and Robert M. Gower. The Polar Express: Optimal matrix sign methods and their application to the Muon algorithm. *arXiv preprint arXiv:2505.16932*, 2025. URL [url]. (cited on pp. 12, 33, 38, and 39)

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 32, 2019. (cited on p. 3)

Arpit Bansal, Avi Schwarzschild, Eitan Borgnia, Zeyad Emam, Furong Huang, Micah Goldblum, and Tom Goldstein. End-to-end algorithm synthesis with recurrent networks: Logical extrapolation without overthinking. In *Advances in Neural Information Processing Systems*, 2022. (cited on pp. 1, 3, and 6)

Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2 edition, 2017. doi: 10.1007/978-3-319-48311-5. (cited on pp. 2, 5, and 22)

Heinz H. Bauschke, Walaa M. Moursi, and Xianfu Wang. Generalized monotone operators and their averaged resolvents. *Mathematical Programming*, 189:55–74, 2021. doi: 10.1007/s10107-020-01500-6. URL [url]. (cited on p. 2)

Jeremy Bernstein and Laker Newhouse. Modular duality in deep learning. *arXiv preprint arXiv:2410.21265*, 2024. (cited on p. 3)

Valérie Castin, Pierre Ablin, and Gabriel Peyré. How smooth is attention? In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 5817–5840, 2024. URL [url]. (cited on pp. 3 and 9)

Franz Louis Cesista. Fast, numerically stable, and auto-differentiable spectral clipping via Newton–Schulz iteration, 2025. URL [url]. (cited on p. 33)

Zizhe Chen, Jiqian Dong, Yizhou Tian, Garry Yang, Yongqiang Chen, Zhitang Chen, and James Cheng. Hista and numca: Estimate state value effectively for large language model reinforcement learning. *arXiv preprint arXiv:2605.29782*, 2026. URL [url]. (cited on p. 8)

A. Chowdhery et al. PaLM: Scaling language modeling with Pathways. *arXiv preprint arXiv:2204.02311*, 2022. (cited on p. 10)

P. L. Combettes and J.-C. Pesquet. Deep neural network structures solving variational inequalities. *Set-Valued and Variational Analysis*, 28:491–518, 2020. (cited on p. 3)

Tri Dao and Albert Gu. Transformers are SSMs: Generalized models and efficient algorithms through structured state space duality. In *Proceedings of the 41st International Conference on Machine Learning*, 2024. URL [url]. (cited on p. 30)

Alec Dewulf, Dhruv Pai, Li Yang, Ashley Zhang, and Ben Keigwin. Aurora: A leverage-aware spectral optimizer, 2026. (cited on pp. 1, 3, 14, and 37)

Jelena Diakonikolas. Pushing the complexity boundaries of fixed-point equations: Adaptation to contraction and controlled expansion. *arXiv preprint arXiv:2506.17698*, 2025. URL [url]. (cited on pp. 2 and 6)

S. M. Emadi. Exact attention sensitivity and the geometry of Transformer stability. *arXiv preprint arXiv:2602.18849*, 2026. (cited on p. 3)

Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach. *arXiv preprint arXiv:2502.05171*, 2025. URL [url]. (cited on pp. 1, 3, 4, and 6)

Henry Gouk, Eibe Frank, Bernhard Pfahringer, and Michael J. Cree. Regularisation of neural networks by enforcing Lipschitz continuity. *Machine Learning*, 110(2):393–416, 2021. doi: 10.1007/s10994-020-05929-w. URL [url]. (cited on p. 3)

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *Conference on Language Modeling*, 2024. URL [url]. (cited on p. 30)

B. Halpern. Fixed points of nonexpanding maps. *Bulletin of the American Mathematical Society*, 73:957–961, 1967. (cited on pp. 1, 2, and 5)

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90. URL [url]. (cited on p. 10)

Nicholas J. Higham. *Functions of Matrices: Theory and Computation*. SIAM, 2008. (cited on pp. 12 and 33)

Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024. URL [url]. (cited on pp. 1 and 33)

Hyunjik Kim, George Papamakarios, and Andriy Mnih. The Lipschitz constant of self-attention. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5562–5571, 2021. URL [url]. (cited on p. 3)

J. Klicpera, A. Bojchevski, and S. Günnemann. Predict then propagate: Graph neural networks meet personalized PageRank. In *International Conference on Learning Representations*, 2019. (cited on p. 3)

Tim Large, Yang Liu, Minyoung Huh, Hyojin Bahng, Phillip Isola, and Jeremy Bernstein. Scalable optimization in the modular norm. In *Advances in Neural Information Processing Systems*, volume 37, 2024. (cited on p. 3)

Tim Tsz-Kit Lau and Weijie Su. Symmetry-compatible principle for optimizer design: Embeddings, LM heads, SwiGLU MLPs, and MoE routers. *arXiv preprint arXiv:2605.18106*, 2026. URL [url]. (cited on p. 3)

Zichong Li, Liming Liu, Chen Liang, Weizhu Chen, and Tuo Zhao. NorMuon: Making Muon more efficient and scalable, 2025. (cited on pp. 1 and 3)

Felix Lieder. On the convergence rate of the Halpern iteration. *Optimization Letters*, 15:405–418, 2021. doi: 10.1007/s11590-020-01617-9. URL [url]. (cited on pp. 2 and 5)

Kai Lion, Florian Hübler, Bingcong Li, Antonio Orvieto, and Niao He. Muown: Row-norm control for Muon optimization, 2026. (cited on pp. 1 and 3)

Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, et al. Muon is scalable for llm training. *arXiv preprint arXiv:2502.16982*, 2025. (cited on p. 31)

Ilya Loshchilov, Cheng-Ping Hsieh, Simeng Sun, and Boris Ginsburg. nGPT: Normalized transformer with representation learning on the hypersphere. In *International Conference on Learning Representations*, 2025. (cited on pp. 1 and 3)

W. R. Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4(3):506–510, 1953. doi: 10.1090/S0002-9939-1953-0054846-3. (cited on p. 2)

Marin Community. Default gate/router weight decay on in the hero recipe, September 2026a. URL [url]. GitHub pull request #8833. (cited on p. 10)

Marin Community. Hero run: Attn gate norm growth, August 2026b. URL [url]. GitHub issue #8818. (cited on p. 10)

Laurent Meunier, Blaise J. Delattre, Alexandre Araujo, and Alexandre Allauzen. A dynamical system perspective for Lipschitz neural networks. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 15484–15500, 2022. URL [url]. (cited on p. 3)

Jisun Park and Ernest K. Ryu. Exact optimal accelerated complexity for fixed-point iterations. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 17420–17457, 2022. URL [url]. (cited on pp. 2 and 5)

Thomas Pethick. When to use what Schatten-p norm in deep learning? *arXiv preprint arXiv:2606.15268*, 2026. (cited on pp. 3 and 12)

Thomas Pethick, Wanyun Xie, and Volkan Cevher. Stable nonconvex-nonconcave training via linear interpolation. *arXiv preprint arXiv:2310.13459*, 2023. URL [url]. (cited on p. 2)

Thomas Pethick, Puya Latafat, Basile Lewandowski, Zheng Xu, Peter Kairouz, Panagiotis Patrinos, and Volkan Cevher. iFedDR: Auto-tuning local computation with inexact Douglas–Rachford splitting in federated learning, 2025a. URL [url]. Preprint. (cited on pp. 12, 32, and 33)

Thomas Pethick, Ioannis Mavrothalassitis, and Volkan Cevher. Efficient interpolation between extragradient and proximal methods for weak mvis. In *International Conference on Learning Representations*, volume 2025, pages 50896–50923, 2025b. (cited on p. 2)

Thomas Pethick, Wanyun Xie, Kimon Antonakopoulos, Zhenyu Zhu, Antonio Silveti-Falls, and Volkan Cevher. Training deep learning models with norm-constrained LMOs. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 49069–49104, 2025c. URL [url]. (cited on pp. 3, 10, and 31)

Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019. doi: 10.1561/2200000073. (cited on p. 27)

Zihan Qiu, Zekun Wang, Bo Zheng, Zeyu Huang, Kaiyue Wen, Songlin Yang, Rui Men, Le Yu, Fei Huang, Suozhi Huang, Dayiheng Liu, Jingren Zhou, and Junyang Lin. Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free. *arXiv preprint arXiv:2505.06708*, 2025. URL [url]. (cited on p. 15)

Qwen Team. Qwen3-Next-80B-A3B: model card and configuration. [url], 2025. (cited on p. 15)

EK Ryu, W Yin, and Large-Scale Convex Optimization. Algorithms & analyses via monotone operators, 2022. (cited on p. 4)

Shoham Sabach and Shimrit Shtern. A first order method for solving convex bilevel optimization problems. *SIAM Journal on Optimization*, 27(2):640–660, 2017. doi: 10.1137/16M105592X. URL [url]. (cited on pp. 2 and 5)

M. E. Sander, P. Ablin, M. Blondel, and G. Peyré. Sinkformers: Transformers with doubly stochastic attention. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 3515–3530, 2022. (cited on pp. 3, 9, and 27)

N. Sato and H. Iiduka. Lipschitz multiscale deep equilibrium models: A theoretically guaranteed and accelerated approach. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, 2026. arXiv:2602.03297. (cited on p. 3)

Noam Shazeer. GLU variants improve Transformer. *arXiv preprint arXiv:2002.05202*, 2020. URL [url]. (cited on pp. 2 and 8)

Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967. doi: 10.2140/pjm.1967.21.343. (cited on p. 27)

S. Sonkar and R. G. Baraniuk. Investigating the role of feed-forward networks in Transformers using parallel attention and feed-forward net design. *arXiv preprint arXiv:2305.13297*, 2023. (cited on p. 10)

Jianlin Su. Computing singular value clipping mclip via msign (part 1), 2025a. URL [url]. Blog post, in Chinese. (cited on p. 33)

Jianlin Su. Computing singular value clipping mclip via msign (part 2), June 2025b. URL [url]. Blog post, in Chinese. (cited on pp. 33 and 34)

Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. In *Advances in Neural Information Processing Systems*, 2018. (cited on p. 35)

Kaiyue Wen, Xingyu Dang, Kaifeng Lyu, Tengyu Ma, and Percy Liang. Fantastic pretraining optimizers and where to find them II: Hyperball optimization, 2026. (cited on pp. 15, 31, and 39)

Ezra Winston and J. Zico Kolter. Monotone operator equilibrium networks. In *Advances in Neural Information Processing Systems*, volume 33, 2020. (cited on p. 3)

R. Wittmann. Approximation of fixed points of nonexpansive mappings. *Archiv der Mathematik*, 58(5):486–491, 1992. (cited on p. 22)

Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention Transformers with hardware-efficient training. In *Proceedings of the 41st International Conference on Machine Learning*, 2024. URL [url]. (cited on p. 30)

Biao Zhang and Rico Sennrich. Root mean square layer normalization. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL [url]. (cited on p. 8)

Contents of the appendix

A Proofs and supplementary results for Halpern iteration **21**
A.1 Preliminary: geometric TBPTT error . . . 21
A.2 Fixed-step Halpern . . . 22
A.3 Logarithmic-step Halpern . . . 22
A.4 Input robustness . . . 23

B Proofs and supplementary results for the neural parameterization **23**
B.1 Normalization profiles and the MLP block . . . 23
B.2 Linearly growing Jacobians, SwiGLU, and squared ReLU . . . 25
B.3 The Sinkhorn-attention block . . . 27
B.4 Gated attention and gated state-space models . . . 28
B.5 Fixed-point existence and uniqueness . . . 30
B.6 Combining attention and MLP blocks . . . 31

C Computing and using the oracles **31**
C.1 Using the LMO inside an optimizer . . . 31
C.2 Sparrow with a row-clipping threshold . . . 31
C.3 Solvers for the joint LMO . . . 32
C.4 Spectral clipping via matrix products . . . 33

D Experimental details **34**
D.1 Robustness-accuracy pareto frontier . . . 34
D.2 Sparrow as a drop-in replacement for Muon . . . 37
D.3 Looped models . . . 40

A Proofs and supplementary results for Halpern iteration

A.1 Preliminary: geometric TBPTT error

Lemma A.1 (Geometric TBPTT error)**.** *Suppose the recurrence is differentiable along its trajectory, u and z^0 are independent of θ, and*

$$A_k := \partial_z \mathbb{G}_{k,\theta,u}(z^k), \qquad B_k := \partial_\theta \mathbb{G}_{k,\theta,u}(z^k)$$

satisfy $\|A_k\| \le q < 1$ and $\|B_k\| \le B$ for every k. For $0 \le H \le K$, let

$$S_K^{(H)} := \textstyle\sum_{i=K-H}^{K-1} A_{K-1} \cdots A_{i+1} B_i,$$

where an empty product is the identity. Then, with $S_K := \partial_\theta z^K$,

$$\|S_K - S_K^{(H)}\| \le \tfrac{B}{1-q} q^H. \tag{18}$$

In particular, a horizon $H = O(\log(1/\varepsilon))$ gives error at most ε, independently of K.

Proof. The chain rule gives $S_{k+1} = A_k S_k + B_k$ and $S_0 = 0$. Iterating,

$$S_K = \textstyle\sum_{i=0}^{K-1} A_{K-1} \cdots A_{i+1} B_i.$$

Therefore

$$\|S_K - S_K^{(H)}\| \le B \textstyle\sum_{i=0}^{K-H-1} q^{K-1-i} = B \sum_{j=H}^{K-1} q^j \le \tfrac{B}{1-q} q^H. \qquad \square$$

A.2 Fixed-step Halpern

Proof of Proposition 4.1. Since $\mathbb{T}_\theta$ is L-Lipschitz,

$$\|\mathbb{H}_{\beta,u}(x) - \mathbb{H}_{\beta,u}(y)\| = (1-\beta)\|\mathbb{T}_\theta(x) - \mathbb{T}_\theta(y)\| \le q\|x-y\|,$$

so $\mathbb{H}_{\beta,u}$ is a contraction with modulus q, and (i) follows from the Banach fixed-point theorem. At every point of differentiability, $\|\partial_z \mathbb{H}_{\beta,u}\| \le q$, so (ii) is Lemma A.1 with this q.

For (iii), rearranging the fixed-point equation $z_\beta = \beta u + (1-\beta)\mathbb{T}_\theta(z_\beta)$ gives

$$z_\beta - \mathbb{T}_\theta(z_\beta) = \tfrac{\beta}{1-\beta}(u - z_\beta), \tag{19}$$

which is the residual identity. Since $\beta > 0$, it shows that $z_\beta \in \mathrm{Fix}(\mathbb{T}_\theta)$ if and only if $z_\beta = u$. Hence, if z_β is a fixed point, then $u = z_\beta \in \mathrm{Fix}(\mathbb{T}_\theta)$. Conversely, if $u \in \mathrm{Fix}(\mathbb{T}_\theta)$, then $\mathbb{H}_{\beta,u}(u) = u$, and uniqueness gives $z_\beta = u$.

Now let $\mathbb{T}_\theta$ be nonexpansive, so that $\mathbb{I} - \mathbb{T}_\theta$ is monotone (Bauschke and Combettes, 2017), and let $c \in \mathcal{C}_\theta$. Since $(\mathbb{I} - \mathbb{T}_\theta)c = 0$, monotonicity and (19) give

$$0 \le \langle (\mathbb{I} - \mathbb{T}_\theta)z_\beta,\, z_\beta - c\rangle = \tfrac{\beta}{1-\beta}\langle u - z_\beta,\, z_\beta - c\rangle,$$

that is, $\langle u - z_\beta, c - z_\beta\rangle \le 0$. Consequently,

$$\|u - c\|^2 = \|u - z_\beta\|^2 + \|z_\beta - c\|^2 + 2\langle u - z_\beta, z_\beta - c\rangle \ge \|u - z_\beta\|^2.$$

Taking the infimum over $c \in \mathcal{C}_\theta$ yields $\|u - z_\beta\| \le \mathrm{dist}(u, \mathcal{C}_\theta)$, and the residual identity gives the stated residual bound. □

A.3 Logarithmic-step Halpern

Proof of Proposition 4.2. The sequence (β_k) lies in $(0,1)$, decreases to zero, and

$$\textstyle\sum_{k=0}^{\infty} \beta_k = \infty, \qquad \sum_{k=0}^{\infty} |\beta_{k+1} - \beta_k| = \beta_0 < \infty.$$

The divergence follows, for example, from $\log(k+3) \le k+3$ and comparison with the harmonic series; the second identity follows by monotonicity and telescoping. The classical Halpern convergence theorem (Wittmann, 1992) therefore proves (i).

For the sensitivity estimates in (ii), write

$$S_k := \partial_\theta z^k, \qquad A_k := (1-\beta_k)\partial_z \mathbb{T}_\theta(z^k), \qquad B_k := (1-\beta_k)\partial_\theta \mathbb{T}_\theta(z^k).$$

Then $S_{k+1} = A_k S_k + B_k$, with

$$\|A_k\| \le 1 - \beta_k, \qquad \|B_k\| \le M.$$

For every $n \ge 1$, iteration of this recursion and monotonicity of (β_k) give

$$\begin{aligned} \|S_n\| &\le M \textstyle\sum_{i=0}^{n-1} \prod_{j=i+1}^{n-1}(1-\beta_j) \\ &\le M \textstyle\sum_{r=0}^{n-1}(1-\beta_{n-1})^r \le \frac{M}{\beta_{n-1}} = M\log(n+2). \end{aligned} \tag{20}$$

The part discarded by H-step TBPTT is exactly

$$S_K - S_K^{(H)} = A_{K-1}\cdots A_{K-H} S_{K-H}.$$

For $H < K$, using $1 - t \le e^{-t}$ and the fact that $\beta_j \ge \beta_{K-1} = 1/\log(K+2)$ when $j \le K-1$, we obtain

$$\begin{aligned} \|S_K - S_K^{(H)}\| &\le \|S_{K-H}\| \textstyle\prod_{j=K-H}^{K-1}(1-\beta_j) \\ &\le M\log(K-H+2)\exp\Big(-\textstyle\sum_{j=K-H}^{K-1}\beta_j\Big) \\ &\le M\log(K+2)\exp\Big(-\tfrac{H}{\log(K+2)}\Big), \end{aligned}$$

which proves (ii). □

Remark A.2 (Bounded full sensitivities). If the full sensitivities also satisfy $\sup_k \|\partial_\theta z^k\| \le S$, replacing (20) by $\|S_{K-H}\| \le S$ in the proof above gives

$$\|S_K - S_K^{(H)}\| \le S \exp\Big(-\tfrac{H}{\log(K+2)}\Big),$$

so a window $H = O(\log K \, \log(S/\varepsilon))$ suffices. This improves on Proposition 4.2(ii) only if S is smaller than the bound $\|\partial_\theta z^k\| \le M \log(k+2)$ that (20) already provides under the assumptions of the proposition.

A.4 Input robustness

Proof of Theorem 4.3. Write $d_k := \|z^k - \widetilde{z}^k\|$. Subtracting the two updates and using that $\mathbb{T}_\theta$ is L-Lipschitz gives

$$d_{k+1} \le \beta_k \|u - \widetilde{u}\| + (1-\beta_k) L \, d_k \le \beta_k \varepsilon + (1-\beta_k) L \, d_k. \tag{21}$$

If $d_k \le C\varepsilon$, then (21) gives

$$d_{k+1} \le \big(\beta_k + (1-\beta_k) LC\big)\varepsilon,$$

which is at most $C\varepsilon$ provided

$$\beta_k \le C\big(1 - (1-\beta_k)L\big). \tag{22}$$

For $L \le 1$ and $C = 1$, (22) reads $(1-\beta_k)L \le 1-\beta_k$, which holds. For $L > 1$, the condition $\beta_k \ge \beta_{\min} > 1 - 1/L$ makes $1 - (1-\beta_k)L > 0$, so (22) reads $f(\beta_k) \le C$ for

$$f(\beta) := \tfrac{\beta}{1-(1-\beta)L}, \qquad f'(\beta) = \tfrac{1-L}{\big(1-(1-\beta)L\big)^2} < 0 \qquad \text{on } (1-1/L, 1].$$

Thus f is decreasing with $f(1) = 1$, so $f(\beta_k) \le f(\beta_{\min}) = C$. In both cases $C \ge 1$, so $d_0 \le \varepsilon \le C\varepsilon$, and induction proves (5).

It remains to show that the constant C cannot be improved, that is, that there is an operator and a schedule for which the separation d_k approaches $C\varepsilon$. Take a fixed schedule $\beta_k \equiv \beta$, so that $\beta_{\min} = \beta$, and $\mathbb{T}_\theta = L\mathbb{I}$ with $L > 1$ and $(1-\beta)L < 1$. Solving the fixed-point equation $z = \beta u + (1-\beta)Lz$ gives

$$z_\beta = \tfrac{\beta}{1-(1-\beta)L}\, u = Cu, \qquad \text{so} \qquad \|z_\beta - \widetilde{z}_\beta\| = C\|u - \widetilde{u}\|,$$

where C is the constant (4) for this schedule. By Proposition 4.1(i), $z^k \to z_\beta$ and $\widetilde{z}^k \to \widetilde{z}_\beta$, so $d_k \to C\|u - \widetilde{u}\|$, which equals $C\varepsilon$ when $\|u - \widetilde{u}\| = \varepsilon$. Hence, for any $C' < C$, the bound $d_k \le C'\varepsilon$ fails for all sufficiently large k, so (5) does not hold with a smaller constant. □

B Proofs and supplementary results for the neural parameterization

B.1 Normalization profiles and the MLP block

The normalizations of Section 5 are instances of the following construction.

Definition B.1 (Normalization profile)**.** *Let $\tau > 0$. A* normalization profile *is a concave, continuously differentiable $h : [0,\infty) \to [0,\tau]$ with $h(0) = 0$ and $h'(0) \le 1$. It defines the radial normalization and constant*

$$\mathbb{N}_h(x) := \tfrac{h(\mathrm{rms}(x))}{\mathrm{rms}(x)}\, x, \quad \mathbb{N}_h(0) := 0, \qquad \omega_h := \sup_{\rho > 0} \tfrac{h(\rho)^2}{\rho}.$$

We call $\mathbb{N}_h$ a normalization layer.

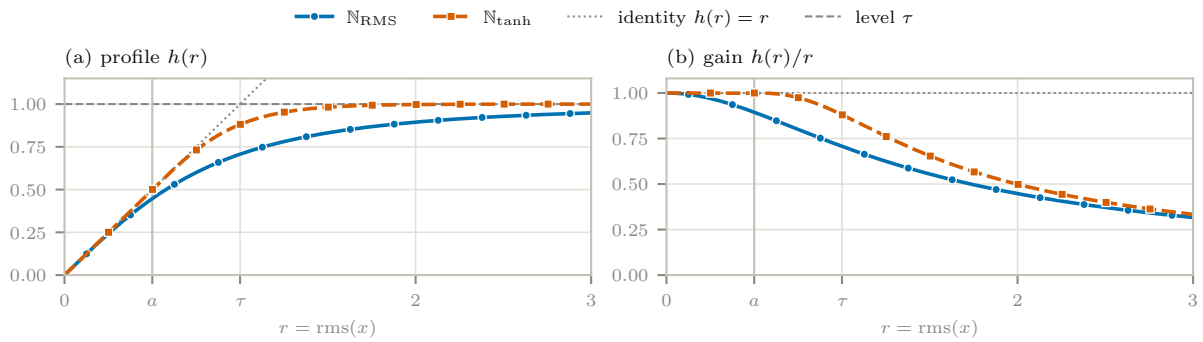

Figure 6: Profiles of the normalization layers $\mathbb{N}_{\mathrm{RMS}}$ and $\mathbb{N}_{\tanh}$ in (6) with $\tau = 1$ and $a = \frac{1}{2}$, as functions of $r = \mathrm{rms}(x)$: (a) the profile $h(r)$ and (b) the gain $h(r)/r = \|D\mathbb{N}_h(x)\|_{\mathrm{op}}$ (24). Both are nonexpansive, since their gains stay at most 1. The gain of $\mathbb{N}_{\tanh}$ equals 1 up to $r = a$ so its compositions with other nonexpansive maps can have multiple fixed points.

Lemma B.2 (Properties of normalization layers)**.** *For every normalization profile h, $\mathbb{N}_h$ is continuously differentiable away from 0, and for all $x, y \in \mathbb{R}^d$:*

(N1) $\|\mathbb{N}_h(x) - \mathbb{N}_h(y)\| \le \|x - y\|$;

(N2) $\|\mathbb{N}_h(x)\| \le \sqrt{d}\,\tau$;

(N3) $\|\mathbb{N}_h(x)\|\,\|D\mathbb{N}_h(x)\|_{\mathrm{op}} \le \sqrt{d}\,\omega_h$ *for* $x \neq 0$, *with* $\omega_h \le \tau$.

Proof. By concavity, h' is nonincreasing, and $h' \ge 0$ since otherwise $h(\rho) \to -\infty$ as $\rho \to \infty$, contradicting $h \ge 0$. Hence $\rho\, h'(\rho) \le h(\rho) = \int_0^\rho h'(s)\, ds \le \rho\, h'(0) \le \rho$, that is,

$$0 \le h'(\rho) \le \tfrac{h(\rho)}{\rho} \le 1 \qquad \text{for every } \rho > 0. \tag{23}$$

Fix $x \neq 0$, and write $\hat{x} := x/\|x\|$ and $\rho := \mathrm{rms}(x)$, so that $\mathbb{N}_h(x) = \sqrt{d}\,h(\rho)\hat{x}$. Differentiating the radial factor and the direction separately gives

$$D\mathbb{N}_h(x) = h'(\rho)\,\hat{x}\hat{x}^\top + \tfrac{h(\rho)}{\rho}\big(\mathbb{I} - \hat{x}\hat{x}^\top\big),$$

which is continuous in $x \neq 0$. This symmetric matrix has eigenvalue $h'(\rho)$ along $\hat{x}$ and $h(\rho)/\rho$ on its orthogonal complement, so (23) gives

$$\|D\mathbb{N}_h(x)\|_{\mathrm{op}} = \tfrac{h(\rho)}{\rho} \le 1. \tag{24}$$

For (N1), $\mathbb{N}_h$ is continuous and a nonconstant line segment meets 0 at most once, so integrating (24) along the segment between x and y gives $\|\mathbb{N}_h(x) - \mathbb{N}_h(y)\| \le \|x - y\|$. For (N2), $\|\mathbb{N}_h(x)\| = \sqrt{d}\,h(\rho) \le \sqrt{d}\,\tau$. For (N3), (24) gives

$$\|\mathbb{N}_h(x)\|\,\|D\mathbb{N}_h(x)\|_{\mathrm{op}} = \sqrt{d}\,\tfrac{h(\rho)^2}{\rho} \le \sqrt{d}\,\omega_h,$$

and $h(\rho)^2/\rho \le h(\rho) \le \tau$ by (23), so $\omega_h \le \tau$. □

Figure 6 shows the profiles of $\mathbb{N}_{\mathrm{RMS}}$ and $\mathbb{N}_{\tanh}$ in (6), $h(r) = \tau r/\sqrt{\tau^2 + r^2}$ and $h(r) = r$ for $r \le a$, $a + (\tau - a)\tanh\big(\frac{r-a}{\tau-a}\big)$ for $r > a$.

The block theorems of Section 5 use only (N1)–(N3), so they hold for every normalization layer, and indeed for any map with these three properties.

Lemma B.3 (The two normalizations)**.** $\mathbb{N}_{\mathrm{RMS}}$ *and* $\mathbb{N}_{\tanh}$ *are normalization layers. They satisfy (N1)–(N3) with* $\omega = \tau/2$ *and* $a \le \omega \le \tau$*, respectively, and* $\mathbb{N}_{\mathrm{RMS}}$ *strictly decreases every nonzero distance.*

Proof. By Lemma B.2, it suffices to write each normalization as $\mathbb{N}_h$, check that h is a normalization profile, and compute $\omega = \omega_h$.

First, $\mathbb{N}_{\mathrm{RMS}} = \mathbb{N}_h$ with

$$h(\rho) = \frac{\tau\rho}{\sqrt{\tau^2+\rho^2}}, \qquad h'(\rho) = \frac{\tau^3}{(\tau^2+\rho^2)^{3/2}}.$$

Thus $h(0) = 0$, $h < \tau$, and h' is continuous and decreases from $h'(0) = 1$, so h is a normalization profile. Its constant is

$$\omega = \sup_{\rho>0} \frac{\tau^2\rho}{\tau^2+\rho^2} = \frac{\tau}{2},$$

attained at $\rho = \tau$. Moreover, $h(\rho)/\rho = \tau/\sqrt{\tau^2+\rho^2} < 1$ for $\rho > 0$, so (24) is strict at every $x \neq 0$. Integrating along the segment between $x \neq y$, as in the proof of (N1), therefore gives

$$\|\mathbb{N}_{\mathrm{RMS}}(x) - \mathbb{N}_{\mathrm{RMS}}(y)\| < \|x - y\|,$$

which is the strict decrease used for the normalization layers in Section 5.

Second, $\mathbb{N}_{\tanh} = \mathbb{N}_h$ with

$$h(\rho) = \min(\rho, a) + (\tau - a)\tanh\Big(\tfrac{(\rho-a)_+}{\tau-a}\Big), \qquad h'(\rho) = \begin{cases} 1, & \rho \le a, \\ \operatorname{sech}^2\big(\frac{\rho-a}{\tau-a}\big), & \rho > a. \end{cases}$$

Thus $h(0) = 0$, $h < \tau$, and h' is continuous and nonincreasing from $h'(0) = 1$, so h is a normalization profile. Since $h(\rho)^2/\rho = \rho$ on $(0, a]$, its constant satisfies $\omega \ge a$, and $\omega \le \tau$ by Lemma B.2. □

Proof of Theorem 5.1. The raw MLP satisfies

$$\|\mathbb{F}_{\mathrm{MLP}}(x) - \mathbb{F}_{\mathrm{MLP}}(\widetilde{x})\| \le \|W_2\|_{\mathrm{op}} L_\sigma \|W_1\|_{\mathrm{op}} \|x - \widetilde{x}\|,$$

where the biases cancel in the difference. Since $\mathbb{N}$ is 1-Lipschitz by (N1), composing it before or after $\mathbb{F}_{\mathrm{MLP}}$ does not increase this bound, which proves (9). □

B.2 Linearly growing Jacobians, SwiGLU, and squared ReLU

Lemma B.4 (Input normalization for linearly growing Jacobians)**.** *Let* $\mathbb{G} : \mathbb{R}^d \to \mathbb{R}^d$ *be continuously differentiable, and suppose that, for some* $L_1, L_0 \ge 0$*,*

$$\|D\mathbb{G}(z)\|_{\mathrm{op}} \le L_1\|z\| + L_0 \qquad \text{for every } z \in \mathbb{R}^d. \tag{25}$$

Then, for every normalization layer $\mathbb{N}$ *(Definition B.1),*

$$\mathrm{Lip}(\mathbb{G} \circ \mathbb{N}) \le L_1\sqrt{d}\,\omega + L_0. \tag{26}$$

For $\mathbb{N}_{\mathrm{RMS}}$*, the bound is* $L_1\sqrt{d}\,\tau/2 + L_0$*.*

Proof. At $x \neq 0$, the chain rule, (N3) and $\|D\mathbb{N}(x)\|_{\mathrm{op}} \le 1$, which follows from (N1), give

$$\|D(\mathbb{G} \circ \mathbb{N})(x)\|_{\mathrm{op}} \le \big(L_1\|\mathbb{N}(x)\| + L_0\big)\|D\mathbb{N}(x)\|_{\mathrm{op}} \le L_1\sqrt{d}\,\omega + L_0.$$

Integrating the Jacobian along line segments, each of which meets 0 at most once, proves (26). □

We are now ready to make the main argument that makes us arrive at a joint spectral and row constraint for matrices followed by ReLU^2 and gated activation functions such as in SwiGLU.

Theorem B.5 (Lipschitz squared-ReLU block)**.** *Let $W_1 : \mathbb{R}^d \to \mathbb{R}^m$ and $W_2 : \mathbb{R}^m \to \mathbb{R}^d$ be linear maps, with $b_1 \in \mathbb{R}^m$ and $b_2 \in \mathbb{R}^d$, let $\mathbb{N}$ be a normalization layer (Definition B.1), and define*

$$\mathbb{M}_{\mathrm{ReLU}^2}(x) := W_2 \,\mathrm{ReLU}(W_1 \mathbb{N}(x) + b_1)^{\odot 2} + b_2. \tag{27}$$

Then

$$\mathrm{Lip}(\mathbb{M}_{\mathrm{ReLU}^2}) \le 2\|W_2\|_{\mathrm{op}}\|W_1\|_{\mathrm{op}} \left[\sqrt{d}\,\omega\|W_1\|_{2\to\infty} + \|\,\mathrm{ReLU}(b_1)\|_\infty\right]. \tag{28}$$

Proof. Define

$$\mathbb{H}(z) := W_2 \,\mathrm{ReLU}(W_1 z + b_1)^{\odot 2} + b_2.$$

Since $t \mapsto \mathrm{ReLU}(t)^2$ has derivative $2\,\mathrm{ReLU}(t)$,

$$D\mathbb{H}(z) = 2W_2 \,\mathrm{diag}(\mathrm{ReLU}(W_1 z + b_1))\, W_1.$$

The operator norm of a diagonal matrix is the vector ℓ_∞ norm over the entries. Moreover, Cauchy–Schwarz applied to each row gives

$$\|W_1 z\|_\infty \le \max_i \|(W_1)_{i,:}\|_2 \,\|z\| = \|W_1\|_{2\to\infty}\|z\|.$$

Using $\mathrm{ReLU}(u+v) \le |u| + \mathrm{ReLU}(v)$ coordinatewise, we therefore obtain

$$\begin{aligned}\|D\mathbb{H}(z)\|_{\mathrm{op}} &\le 2\|W_2\|_{\mathrm{op}}\|W_1\|_{\mathrm{op}}\|\,\mathrm{ReLU}(W_1 z + b_1)\|_\infty \\ &\le 2\|W_2\|_{\mathrm{op}}\|W_1\|_{\mathrm{op}}\left(\|W_1\|_{2\to\infty}\|z\| + \|\,\mathrm{ReLU}(b_1)\|_\infty\right).\end{aligned}$$

Thus its Jacobian has the linear-growth structure of Lemma B.4. Applying Lemma B.4 with

$$L_1 = 2\|W_2\|_{\mathrm{op}}\|W_1\|_{\mathrm{op}}\|W_1\|_{2\to\infty}, \qquad L_0 = 2\|W_2\|_{\mathrm{op}}\|W_1\|_{\mathrm{op}}\|\,\mathrm{ReLU}(b_1)\|_\infty$$

proves (28). □

Proof of Theorem 5.3. Let $z = \mathbb{N}(x)$ and $\mathbb{G}(z) := W_o[\mathrm{SiLU}(W_g z) \odot W_v z]$. The product rule gives

$$D\mathbb{G}(z) = W_o\left[\mathrm{Diag}\big(\mathrm{SiLU}'(W_g z) \odot W_v z\big)\, W_g + \mathrm{Diag}(\mathrm{SiLU}(W_g z))\, W_v\right].$$

Since $|\,\mathrm{SiLU}(t)| \le |t|$ and $|\,\mathrm{SiLU}'(t)| \le L_{\mathrm{SiLU}}$, the row-norm bounds $\|W_g z\|_\infty \le \|W_g\|_{2\to\infty}\|z\|$ and $\|W_v z\|_\infty \le \|W_v\|_{2\to\infty}\|z\|$ imply

$$\|D\mathbb{G}(z)\|_{\mathrm{op}} \le C\|z\|, \quad C := \|W_o\|_{\mathrm{op}}\left[L_{\mathrm{SiLU}}\|W_v\|_{2\to\infty}\|W_g\|_{\mathrm{op}} + \|W_g\|_{2\to\infty}\|W_v\|_{\mathrm{op}}\right]. \tag{29}$$

Thus (25) in Lemma B.4 holds with $L_1 = C$ and $L_0 = 0$. Lemma B.4 gives (11). □

Remark B.6 (Alternative SwiGLU certificates). Applied to a direction h, the derivative in the proof above is

$$D\mathbb{G}(z)h = W_o\big[(W_v z) \odot \big(\mathrm{SiLU}'(W_g z) \odot W_g h\big) + \mathrm{SiLU}(W_g z) \odot (W_v h)\big].$$

Each product $a \odot b$ can be bounded by $\|a\|_\infty\|b\|$ or by $\|a\|\|b\|_\infty$, and the ℓ_∞ norm of a projection is controlled by its row norm. With $r_\cdot := \|W_\cdot\|_{2\to\infty}$ and $s_\cdot := \|W_\cdot\|_{\mathrm{op}}$, the four choices give the constants $C/\|W_o\|_{\mathrm{op}}$ below, where (i) is the one used in Theorem 5.3.

	ℓ_∞ in first term	ℓ_∞ in second term	$C/\|W_o\|_{\mathrm{op}}$	Norms used
(i)	$W_v z$	$\mathrm{SiLU}(W_g z)$	$L_{\mathrm{SiLU}} r_v s_g + r_g s_v$	rows and spectra of both
(ii)	$W_g h$	$\mathrm{SiLU}(W_g z)$	$(1 + L_{\mathrm{SiLU}}) r_g s_v$	row of W_g, spectrum of W_v
(iii)	$W_v z$	$W_v h$	$(1 + L_{\mathrm{SiLU}}) r_v s_g$	row of W_v, spectrum of W_g
(iv)	$W_g h$	$W_v h$	$L_{\mathrm{SiLU}} r_g s_v + r_v s_g$	rows and spectra of both

The four coincide when $r_v/s_v = r_g/s_g$, in particular for matrices saturated in both norms, where each equals $(1 + L_{\mathrm{SiLU}})rs$. This is the case used in Remark 5.4 to choose the predetermined radii of Sparrow.

B.3 The Sinkhorn-attention block

We use Sinkhorn based attention, following Sinkformers (Sander et al., 2022), since its doubly stochastic attention matrices are nonexpansive. Equip $\mathbb{R}^{n\times d}$ with the Frobenius norm. Write $\mathcal{B}_n := \{P \in \mathbb{R}_+^{n\times n} : P\mathbf{1} = \mathbf{1},\ P^\top\mathbf{1} = \mathbf{1}\}$ for the Birkhoff polytope, and define the exact entropy-regularized Sinkhorn map by

$$\mathrm{Sink}(C) := \arg\max_{P\in\mathcal{B}_n} \left\{ \langle C, P\rangle_F - \textstyle\sum_{i,j} P_{ij}\log P_{ij} \right\}. \tag{30}$$

Theorem B.7 (Lipschitz Sinkhorn-attention block)**.** *Let $W_Q, W_K : \mathbb{R}^d \to \mathbb{R}^{d_k}$, $W_V : \mathbb{R}^d \to \mathbb{R}^{d_v}$, and $W_O : \mathbb{R}^{d_v} \to \mathbb{R}^d$ be linear maps. Fix a normalization layer $\mathbb{N}$ (Definition B.1), and positive spectral upper bounds $\kappa_Q \geq \|W_Q\|_{\mathrm{op}}$ and $\kappa_K \geq \|W_K\|_{\mathrm{op}}$. Define*

$$\mathbb{Z}(X) := \mathbb{N}(X), \qquad \mathbb{Q}(X) := \mathbb{Z}(X)W_Q, \qquad \mathbb{K}(X) := \mathbb{Z}(X)W_K, \qquad \mathbb{V}(X) := \mathbb{Z}(X)W_V. \tag{31}$$

For an inverse-temperature parameter $B > 0$, set

$$\mathbb{C}(X) := \tfrac{B\,\mathbb{Q}(X)\mathbb{K}(X)^\top}{d\tau^2\kappa_Q\kappa_K}, \qquad \mathbb{W}(X) := \mathrm{Sink}(\mathbb{C}(X)). \tag{32}$$

Then $|[\mathbb{C}(X)]_{ij}| \leq B$. Define the attention block by

$$\mathbb{A}(X) := \mathbb{W}(X)\mathbb{V}(X)W_O. \tag{33}$$

Then, in the Frobenius norm,

$$\mathrm{Lip}(\mathbb{A}) \leq \left[1 + Be^{4B}\left(\tfrac{\|W_Q\|_{\mathrm{op}}}{\kappa_Q} + \tfrac{\|W_K\|_{\mathrm{op}}}{\kappa_K}\right)\right]\|W_V\|_{\mathrm{op}}\|W_O\|_{\mathrm{op}} \leq (1+2Be^{4B})\|W_V\|_{\mathrm{op}}\|W_O\|_{\mathrm{op}}, \tag{34}$$

independently of the sequence length n.

Proof of Theorem B.7. We first record the sensitivity of exact Sinkhorn normalization on the bounded score range. Let $P = \mathrm{Sink}(C)$ with $|C_{ij}| \leq B$. Since P is a diagonal scaling of $\exp(C)$ (Sinkhorn and Knopp, 1967; Peyré and Cuturi, 2019), write $P_{ij} = a_i e^{C_{ij}} b_j$. The column constraints give

$$\tfrac{b_j}{b_k} = \frac{\sum_i a_i e^{C_{ik}}}{\sum_i a_i e^{C_{ij}}} \leq e^{2B},$$

hence $P_{ij}/P_{ik} \leq e^{4B}$. Because each row sums to one,

$$0 < P_{ij} \leq \rho_n(B) := \tfrac{e^{4B}}{n-1+e^{4B}}. \tag{35}$$

Now let $\widetilde{P} = \mathrm{Sink}(\widetilde{C})$ with $|\widetilde{C}_{ij}| \leq B$, and let $\Phi(P) := \sum_{i,j} P_{ij}(\log P_{ij} - 1)$, which differs from the regularizer in (30) by the constant n on $\mathcal{B}_n$. The optimality conditions for (30) at P and $\widetilde{P}$, tested with $\widetilde{P}$ and P respectively, give

$$\langle C - \nabla\Phi(P), \widetilde{P} - P\rangle \leq 0, \qquad \langle \widetilde{C} - \nabla\Phi(\widetilde{P}), P - \widetilde{P}\rangle \leq 0.$$

By (35), the Hessian $\nabla^2\Phi = \mathrm{diag}(1/P_{ij})$ is at least $\rho_n(B)^{-1}\mathbb{I}$ on the segment between P and $\widetilde{P}$. Adding the two conditions and using Cauchy–Schwarz therefore gives

$$\rho_n(B)^{-1}\|P - \widetilde{P}\|_F^2 \leq \langle \nabla\Phi(P) - \nabla\Phi(\widetilde{P}), P - \widetilde{P}\rangle \leq \langle C - \widetilde{C}, P - \widetilde{P}\rangle \leq \|C - \widetilde{C}\|_F\|P - \widetilde{P}\|_F,$$

that is,

$$\|\,\mathrm{Sink}(C) - \mathrm{Sink}(\widetilde{C})\|_F \leq \rho_n(B)\|C - \widetilde{C}\|_F. \tag{36}$$

Write $\Delta := \|\mathbb{Z}(X) - \mathbb{Z}(Y)\|_F$. By (N1) and (N2) for the shared input normalization, applied rowwise, $\Delta \le \|X - Y\|_F$ and

$$\begin{aligned} & & \|z_i(X)\| &\le \sqrt{d}\,\tau, \\ \|\mathbb{Q}(X) - \mathbb{Q}(Y)\|_F &\le \|W_Q\|_{\mathrm{op}}\Delta, & \|q_i(X)\| &\le \sqrt{d}\,\tau\kappa_Q, \\ \|\mathbb{K}(X) - \mathbb{K}(Y)\|_F &\le \|W_K\|_{\mathrm{op}}\Delta, & \|k_i(X)\| &\le \sqrt{d}\,\tau\kappa_K, \\ \|\mathbb{V}(X) - \mathbb{V}(Y)\|_F &\le \|W_V\|_{\mathrm{op}}\Delta, & \|v_i(X)\| &\le \sqrt{d}\,\tau\|W_V\|_{\mathrm{op}}. \end{aligned} \tag{37}$$

The row bounds give $\|\mathbb{Q}(X)\|_{\mathrm{op}} \le \|\mathbb{Q}(X)\|_F \le \sqrt{nd}\,\tau\kappa_Q$, and likewise for $\mathbb{K}$ and $\mathbb{V}$. Splitting $\mathbb{Q}(X)\mathbb{K}(X)^\top - \mathbb{Q}(Y)\mathbb{K}(Y)^\top$ into a difference in $\mathbb{Q}$ and a difference in $\mathbb{K}$ and dividing by $d\tau^2\kappa_Q\kappa_K$ then gives

$$\|\mathbb{C}(X) - \mathbb{C}(Y)\|_F \le \tfrac{B\sqrt{n}}{\sqrt{d}\,\tau}\left(\tfrac{\|W_Q\|_{\mathrm{op}}}{\kappa_Q} + \tfrac{\|W_K\|_{\mathrm{op}}}{\kappa_K}\right)\Delta. \tag{38}$$

For the attention output, split

$$\mathbb{W}(X)\mathbb{V}(X) - \mathbb{W}(Y)\mathbb{V}(Y) = \mathbb{W}(X)\big(\mathbb{V}(X) - \mathbb{V}(Y)\big) + \big(\mathbb{W}(X) - \mathbb{W}(Y)\big)\mathbb{V}(Y).$$

Since $\mathbb{W}(X)$ is doubly stochastic, it is a convex combination of permutation matrices by the Birkhoff–von Neumann theorem, so $\|\mathbb{W}(X)\|_{\mathrm{op}} = 1$ and the first term is at most $\|W_V\|_{\mathrm{op}}\Delta$ in the Frobenius norm. By (36), (38) and $\|\mathbb{V}(Y)\|_{\mathrm{op}} \le \sqrt{nd}\,\tau\|W_V\|_{\mathrm{op}}$, the second is at most $Bn\rho_n(B)\big(\tfrac{\|W_Q\|_{\mathrm{op}}}{\kappa_Q} + \tfrac{\|W_K\|_{\mathrm{op}}}{\kappa_K}\big)\|W_V\|_{\mathrm{op}}\Delta$. Finally, $n\rho_n(B) = \frac{ne^{4B}}{n-1+e^{4B}} \le e^{4B}$, so applying W_O gives

$$\|\mathbb{A}(X) - \mathbb{A}(Y)\|_F \le \left[1 + Be^{4B}\left(\tfrac{\|W_Q\|_{\mathrm{op}}}{\kappa_Q} + \tfrac{\|W_K\|_{\mathrm{op}}}{\kappa_K}\right)\right]\|W_V\|_{\mathrm{op}}\|W_O\|_{\mathrm{op}}\|\mathbb{N}(X) - \mathbb{N}(Y)\|_F. \tag{39}$$

Since $\|\mathbb{N}(X) - \mathbb{N}(Y)\|_F = \Delta \le \|X - Y\|_F$, this proves the first bound in (34); the bounds $\|W_Q\|_{\mathrm{op}} \le \kappa_Q$ and $\|W_K\|_{\mathrm{op}} \le \kappa_K$ give the second.

□

B.4 Gated attention and gated state-space models

Gated attention and gated SSMs both multiply a bounded signal elementwise by a sigmoid gate computed from the normalized input. Similarly to the ReLU^2 based MLP and the SwiGLU block (Theorems B.5 and 5.3), the gate acts as a diagonal matrix, whose operator norm is its largest entry, so its sensitivity ends up being controlled by the row norms of the gate projection.

Lemma B.8 (Sigmoid gates)**.** *Let $\alpha(z) := \sigma(Wz + b)$ with $\sigma(t) := (1 + e^{-t})^{-1}$. Then*

$$\|\operatorname{Diag}(\alpha(z) - \alpha(\widetilde{z}))\|_{\mathrm{op}} = \|\alpha(z) - \alpha(\widetilde{z})\|_\infty \le \tfrac{1}{4}\|W\|_{2\to\infty}\|z - \widetilde{z}\|, \tag{40}$$

and the bound is unchanged if each gate is repeated across several channels.

Proof. Since $|\sigma'| \le 1/4$, the ith gate changes by at most $\frac{1}{4}|\langle W_{i,:}, z - \widetilde{z}\rangle| \le \frac{1}{4}\|W_{i,:}\|\|z - \widetilde{z}\|$. Repeating entries does not change the maximum. □

Both certificates below split a gated product as

$$D(x)u(x) - D(y)u(y) = D(x)\big(u(x) - u(y)\big) + \big(D(x) - D(y)\big)u(y), \tag{41}$$

where D is the diagonal gate matrix. The first term is bounded by the gate magnitude, and the second by Lemma B.8 times a bound M on $\|u\|$, which comes from the attention heads in gated attention and from the stability of the recurrence in a gated SSM. Both blocks are also continuous functions of the normalized input for each fixed sequence length (for attention, with the Sinkhorn heads of Theorem B.7), so Theorem B.11 guarantees a fixed point regardless of the Lipschitz bound.

B.4.1 Gated attention

Theorem B.9 (Lipschitz headwise-gated attention)**.** *Let $\mathbb{N}$ be a normalization layer (Definition B.1) applied to each row, and let $\mathbb{U}(X) = [\mathbb{U}_1(X) \mid \cdots \mid \mathbb{U}_H(X)] \in \mathbb{R}^{n\times Hd_v}$ be concatenated attention-head outputs with $\mathrm{Lip}(\mathbb{U}) \le L_U$ in the Frobenius norm and $\sup_{X,i} \|\mathbb{U}(X)_{i,:}\| \le M$. For a gate projection $W_a \in \mathbb{R}^{H\times d}$, $b \in \mathbb{R}^H$, $c > 0$, and an output projection $W_O \in \mathbb{R}^{Hd_v\times d}$, define*

$$g(X) := c\,\sigma\Big(\mathbb{N}(X)W_a^\top + \mathbf{1}b^\top\Big), \qquad \mathbb{A}_g(X) := [g_{:,1}(X)\odot\mathbb{U}_1(X) \mid \cdots \mid g_{:,H}(X)\odot\mathbb{U}_H(X)]W_O, \tag{42}$$

where each gate is broadcast across its head's channels. Then

$$\mathrm{Lip}(\mathbb{A}_g) \le c\|W_O\|_{\mathrm{op}}\left(L_U + \tfrac{M}{4}\|W_a\|_{2\to\infty}\right). \tag{43}$$

Proof. At token i, apply (41) to the concatenated head output u_i and the diagonal matrix D_i of its broadcast gates. The first term has norm at most $c\|u_i(X) - u_i(Y)\|$, and by Lemma B.8 and (N1) the second has norm at most $\frac{c}{4}M\|W_a\|_{2\to\infty}\|x_i - y_i\|$. Taking the Euclidean norm across tokens, using the triangle inequality, and applying W_O proves the bound. □

For the Sinkhorn heads of Theorem B.7, $L_U \le (1 + 2Be^{4B})S_V$ and $M \le \sqrt{d}\,\tau S_V$ with $S_V := (\sum_h \|W_{V,h}\|_{\mathrm{op}}^2)^{1/2}$, so the bound is independent of the sequence length.

B.4.2 Gated state-space models

In a gated SSM the gate multiplies the previous state, so the bound M on the gated signal comes from the stability of the recurrence.

Theorem B.10 (Lipschitz gated state-space models)**.** *Let $\mathbb{N}$ be a normalization layer (Definition B.1) on $\mathbb{R}^d$, $W_a, W_V \in \mathbb{R}^{m\times d}$, $b_a \in \mathbb{R}^m$, and $W_O \in \mathbb{R}^{d\times m}$. Define $\mathbb{S}_g(X) = (y^1, \ldots, y^n)$ by*

$$z^k := \mathbb{N}(x^k), \qquad \alpha(z) := \sigma(W_a z + b_a), \qquad h^k := \alpha(z^k)\odot h^{k-1} + W_V z^k, \qquad y^k := W_O h^k, \tag{44}$$

with $h^0 = 0$. Set

$$\rho := \max_i \sigma\Big((b_a)_i + \sqrt{d}\,\tau\|(W_a)_{i,:}\|\Big) < 1. \tag{45}$$

Then, in the Frobenius norm,

$$\mathrm{Lip}(\mathbb{S}_g) \le \tfrac{\|W_O\|_{\mathrm{op}}\|W_V\|_{\mathrm{op}}}{1-\rho}\left(1 + \tfrac{\sqrt{d}\,\tau\|W_a\|_{2\to\infty}}{4(1-\rho)}\right), \tag{46}$$

independently of the sequence length n.

Proof. By (N2), $\|z^k\| \le \sqrt{d}\,\tau$, so $\|\alpha(z^k)\|_\infty \le \rho$ and

$$\|h^k\| \le \rho\|h^{k-1}\| + \sqrt{d}\,\tau\|W_V\|_{\mathrm{op}}.$$

Since $h^0 = 0$, induction gives $\|h^k\| \le M := \sqrt{d}\,\tau\|W_V\|_{\mathrm{op}}/(1-\rho)$.

For two inputs, let $e^k := \|h^k - \widetilde{h}^k\|$ and $\delta^k := \|x^k - \widetilde{x}^k\|$. Applying (41) to $\alpha(z^k)\odot h^{k-1}$,

$$h^k - \widetilde{h}^k = \alpha(z^k)\odot\big(h^{k-1} - \widetilde{h}^{k-1}\big) + \big(\alpha(z^k) - \alpha(\widetilde{z}^k)\big)\odot\widetilde{h}^{k-1} + W_V\big(z^k - \widetilde{z}^k\big).$$

The three terms are bounded by ρe^{k-1}, by $\frac{M}{4}\|W_a\|_{2\to\infty}\delta^k$ using Lemma B.8 and (N1), and by $\|W_V\|_{\mathrm{op}}\delta^k$ using (N1). Hence

$$e^k \le \rho e^{k-1} + K\delta^k, \qquad K := \|W_V\|_{\mathrm{op}} + \tfrac{M}{4}\|W_a\|_{2\to\infty}. \tag{47}$$

Since $e^0 = 0$, iterating (47) gives $e^k \le K\sum_{j\le k}\rho^{k-j}\delta^j$, the convolution of $(K\delta^k)$ with $(1, \rho, \rho^2, \ldots)$, whose ℓ_1 norm is $1/(1-\rho)$. Young's convolution inequality then yields

$$\|\mathbb{S}_g(X) - \mathbb{S}_g(\widetilde{X})\|_F \le \|W_O\|_{\mathrm{op}}\|(e^k)\|_2 \le \tfrac{\|W_O\|_{\mathrm{op}}K}{1-\rho}\|X - \widetilde{X}\|_F,$$

and substituting M and K proves the claim. □

Gate constraint and stability margin To certify a Lipschitz constant L, choose a row-norm radius $R \geq 0$ and a gate cap $\bar{\rho} \in (0, 1)$ such that

$$\begin{gathered} \|W_a\|_{2\to\infty} \leq R, \qquad \max_i (b_a)_i + \sqrt{d}\,\tau R \leq \log \tfrac{\bar{\rho}}{1-\bar{\rho}}, \\ \tfrac{\|W_O\|_{\mathrm{op}}\|W_V\|_{\mathrm{op}}}{1-\bar{\rho}} \left(1 + \tfrac{\sqrt{d}\,\tau R}{4(1-\bar{\rho})}\right) \leq L. \end{gathered} \tag{48}$$

The bias condition ensures $\rho \leq \bar{\rho}$, so the certificate holds uniformly over these constraints. As in gated attention, W_a needs only a row-norm constraint, which here controls both the gate sensitivity and the stability margin.

Input-dependent writes and readouts The same proof applies to $h^k = \alpha(z^k) \odot h^{k-1} + \mathbb{B}(z^k)$ and $y^k = \mathbb{C}(z^k)h^k$, provided $\mathbb{B}$ and $\mathbb{C}$ are bounded and Lipschitz on the normalized-input ball. If $\|\mathbb{B}\| \leq M_B$, the state bound becomes $M = M_B/(1-\rho)$, and the constant picks up the norm and Lipschitz bounds of $\mathbb{B}$ and $\mathbb{C}$. This is the structure of selective state-space models such as Mamba (Gu and Dao, 2024), whose write and readout depend on the input, and of matrix memories $S^k = \mathrm{Diag}(\alpha(z^k))S^{k-1} + k(z^k)v(z^k)^\top$ read out as $W_O(S^k)^\top q(z^k)$, as in gated linear attention (Yang et al., 2024) and, with one gate per head, Mamba-2 (Dao and Gu, 2024). Vectorizing S^k preserves its norm, and repeating a gate preserves $\|W_a\|_{2\to\infty}$. These models use other gate nonlinearities than the sigmoid, which replaces the constant $1/4$ in Lemma B.8 by a bound on the derivative of the gate, and ρ by the largest gate value on the normalized-input ball.

B.5 Fixed-point existence and uniqueness

Normalization layers provide a way of ensuring existence, as the following theorem makes precise.

Theorem B.11 (Fixed points from normalization layer)**.** *Let $\mathbb{F} : \mathcal{H} \to \mathcal{E}$ and $\mathbb{G} : \mathcal{E} \to \mathcal{H}$ be continuous maps between finite-dimensional Hilbert spaces, and let $\mathbb{N} : \mathcal{E} \to \mathcal{E}$ be a normalization layer (Definition B.1). Its bounded-range property (N2) ensures that $\mathbb{T} := \mathbb{G} \circ \mathbb{N} \circ \mathbb{F}$ has a fixed point.*

Theorem B.11 covers the preceding blocks through the decompositions

$$\begin{aligned} \text{MLP:} \quad & \mathbb{M}_{\text{in}} = \mathbb{F}_{\text{MLP}} \circ \mathbb{N} \circ \mathbb{I}, \quad \mathbb{M}_{\text{out}} = \mathbb{I} \circ \mathbb{N} \circ \mathbb{F}_{\text{MLP}}, \\ \text{SwiGLU/ReLU}^2\text{:} \quad & \mathbb{M}_{\text{SwiGLU}} = \mathbb{G}_{\text{SwiGLU}} \circ \mathbb{N} \circ \mathbb{I}, \\ & \mathbb{M}_{\text{ReLU}^2} = \mathbb{G}_{\text{ReLU}^2} \circ \mathbb{N} \circ \mathbb{I}, \\ \text{attention:} \quad & \mathbb{A} = \mathbb{G}_{\text{attn}} \circ \mathbb{N} \circ \mathbb{I}. \end{aligned}$$

Here the $\mathbb{G}$ maps denote the corresponding computations after the normalized input in (10), (27), and (31). All these module families therefore have a fixed point, regardless of their Lipschitz constants.

Proof of Theorem B.11. By (N2), choose $\rho < \infty$ with $\|\mathbb{N}(z)\| \leq \rho$ for all $z \in \mathcal{E}$. The closed ball $\{z \in \mathcal{E} : \|z\| \leq \rho\}$ is compact, so continuity of $\mathbb{G}$ gives

$$R := \max_{\|z\| \leq \rho} \|\mathbb{G}(z)\| < \infty.$$

It follows that $\|\mathbb{T}(x)\| \leq R$ for every $x \in \mathcal{H}$. Therefore $\mathbb{T}$ maps the closed ball of radius R in $\mathcal{H}$ into itself, and Brouwer's fixed-point theorem gives $\mathrm{Fix}(\mathbb{T}) \neq \varnothing$. □

B.6 Combining attention and MLP blocks

Proof of Proposition 5.5. For every $X, Y \in \mathcal{H}$, the triangle inequality gives

$$\|\mathbb{F}_A(X) - \mathbb{F}_A(Y)\| \le (1-\alpha_A)\|X-Y\| + \alpha_A\|\mathbb{A}_\theta(X) - \mathbb{A}_\theta(Y)\| \le (1 - \alpha_A + \alpha_A L_A)\|X-Y\|,$$

and likewise for $\mathbb{F}_M$. Lipschitz constants multiply under composition, which gives the bound for $\mathbb{T}_{\text{ser}}$. For the parallel map,

$$\begin{aligned}\|\mathbb{T}_{\text{par}}(X) - \mathbb{T}_{\text{par}}(Y)\| &\le (1-\alpha)\|\mathbb{A}_\theta(X) - \mathbb{A}_\theta(Y)\| + \alpha\|\mathbb{M}_\theta(X) - \mathbb{M}_\theta(Y)\| \\ &\le \big((1-\alpha)L_A + \alpha L_M\big)\|X-Y\|.\end{aligned}$$

With $L_A, L_M \le 1$, both bounds are at most 1. □

C Computing and using the oracles

C.1 Using the LMO inside an optimizer

The LMO only provides the direction of the update. In practice we combine it with a Nesterov-type momentum estimator, weight decay and an optional weight normalization N:

$$\begin{aligned} m^{k+1} &= \mu_k m^k + (1-\mu_k)\nabla f(x^k, \xi_k) \\ \bar{m}^{k+1} &= \bar{\mu}_k m^{k+1} + (1-\bar{\mu}_k)\nabla f(x^k, \xi_k) \\ x^{k+1} &= \mathsf{N}\big((1-\lambda_k\gamma_k)x^k - \gamma_k \mathsf{lmo}_{\mathcal{K}}(\bar{m}^{k+1})\big)\end{aligned} \tag{49}$$

Here $\nabla f(x^k, \xi_k)$ is a stochastic gradient, $\gamma_k > 0$ is the step size and $\lambda_k \ge 0$ the weight decay, and $\bar{\mu}_k = 1$ gives standard Heavy ball momentum. The map N is either the identity or the rescaling to the initial Frobenius norm,

$$\mathsf{N}(X) := \|x^0\|_F \tfrac{X}{\|X\|_F}, \tag{50}$$

as in Hyperball (Wen et al., 2026), but without normalizing the update direction. For a spectral ball $\mathcal{K}$ this is Muon with decoupled weight decay as studied in Pethick et al. (2025c); Liu et al. (2025), and for $\mathcal{K} = \mathcal{K}_{r,s}$, with the LMO approximated by Algorithm 1, it is Sparrow.

C.2 Sparrow with a row-clipping threshold

Recall from Section 6.1 that one step of forward-backward splitting (FB) on the dual (17), started from $A = \mathsf{RowShrink}_c(G)$, gives $(\mathsf{RowClip}_r \circ \mathsf{Spectral}_s \circ \mathsf{RowClip}_c)(G)$, and that Sparrow is obtained for $c \to 0$ and U_1 for $c \to \infty$. Here we consider arbitrary c.

Algorithm 2 replaces Row_r on the SwiGLU gate and value matrices in Algorithm 1 by $\mathsf{RowClip}_c$. The threshold c is in the units of G, whose scale changes with the gradient size. We therefore set it relative to the root-mean-square row norm of G,

$$\bar{g} := \Big(\tfrac{1}{m}\textstyle\sum_{i=1}^m \|G_{i,:}\|_2^2\Big)^{1/2} = \tfrac{\|G\|_F}{\sqrt{m}}, \qquad c = \tfrac{\rho}{1-\rho}\,\bar{g},$$

with a unitless $\rho \in [0,1]$. Since $\mathsf{Spectral}_s$ is invariant to positive scaling, we can replace $\mathsf{RowClip}_c(G)$ by

$$\tilde{G}_{i,:} = \frac{G_{i,:}}{\max\{(1-\rho)\|G_{i,:}\|_2,\ \rho\bar{g}\}}. \tag{51}$$

We make this change for numerical stability, since this form is defined on all of $[0,1]$ and attains both ends exactly:

- $\rho = 0$: $\tilde{G} = \mathsf{Row}_1(G)$, which gives Sparrow ($U_3$).

- $\rho = \frac{1}{2}$: $c = \bar{g}$, so only the rows longer than the root-mean-square row norm are clipped.
- $\rho = 1$: $\tilde{G} = G/\bar{g}$, so no row is clipped, which gives U_1 in (16).

It also keeps the input of $\mathsf{Spectral}_s$ at a scale that does not depend on G. Like Sparrow, Algorithm 2 applies $\mathsf{Spectral}_s$ at most once and its output lies in $\mathcal{K}_{r,s}$. The parameter ρ need not be fixed during training; one could, for instance, schedule it inversely proportional to the step-size schedule γ_k, so that the update is close to Sparrow at the peak step size and moves towards U_1 as the step size decays.

Algorithm 2 $\mathsf{Sparrow}_\rho$ update direction. For $\rho = 0$ it is Sparrow and for $\rho = 1$ it is U_1.

Require: gradient estimate $G \in \mathbb{R}^{m\times d}$, radii s and $r = s\min\{1, \sqrt{d/m}\}$, $\rho \in [0, 1]$
Return: depending on the (hidden) parameter group of G:

1: SwiGLU gate and value matrices ($m > d$):	$(\mathsf{RowClip}_r \circ \mathsf{Spectral}_s)(\tilde{G})$ with $\tilde{G}$ in (51)
2: Attention-gate/SSM-gate matrices:	$\mathsf{Row}_r(G)$
3: All other matrices:	$\mathsf{Spectral}_s(G)$

C.3 Solvers for the joint LMO

If our coarse approximation used in Sparrow are not sufficient, one can run an iterative solver like DRS to solve for the joint constraint. DRS relies on $\mathsf{SpectralClip}_s$ which is typically only computed approximately. We therefore compare Douglas–Rachford splitting (DRS) with a fixed approximation rank against the inexact variant iDRS (Pethick et al., 2025a), which adapts the rank to the needed accuracy. The problem is the LMO over $\mathcal{K}_{r,s}$, that is, $(a, H) = (0, G)$ in Section 6. We find that iDRS certifies every instance without tuning the rank and is faster than exact DRS throughout (Table 3).

Problem instances The gradients are $G = DU\,\mathrm{diag}(j^{-\nu})V^\top \in \mathbb{R}^{1024\times 256}$ with random orthonormal U, V, row scales D spaced geometrically in $[0.2, 5]$, and $\|G\|_F = 1$. We take $\nu \in \{1, 1.5\}$, $s = 1$, $s/r \in \{6, 8, 12\}$ and three seeds, which gives 18 instances, all with both constraints binding.

Douglas–Rachford splitting For $\min_W f(W) + g(W)$ with convex f and g, DRS iterates, for any $\gamma > 0$,

$$W := \mathrm{prox}_{\gamma g}(Z), \qquad Z \leftarrow Z + \mathrm{prox}_{\gamma f}(2W - Z) - W, \tag{DRS}$$

and W converges to a minimizer. With f and g from Section 6 and $\gamma = 1$, $\mathrm{prox}_{\gamma f}(Z) = \mathsf{RowClip}_r(Z + \gamma G)$ is exact, and $\mathrm{prox}_{\gamma g} = \mathsf{SpectralClip}_s$ is the step we approximate.

Approximate spectral clip We clip only the top k singular values, computed on a subspace $B \in \mathbb{R}^{d\times k}$ that is warm-started across iterations, and rescale the remainder so that the output stays feasible:

$$\begin{aligned} &Q = \mathrm{qr}(ZB), \qquad Q^\top Z = \hat{U}\,\mathrm{diag}(\alpha)\hat{V}^\top, \qquad P = Q\hat{U}, \\ &T = Z(I - \hat{V}\hat{V}^\top), \qquad c = \min\{1, s/\|T\|_F\}, \\ &w = P\,\mathrm{diag}(\min\{\alpha, s\})\hat{V}^\top + cT, \qquad n = \tfrac{1}{\gamma}P\,\mathrm{diag}((\alpha - s)_+)\hat{V}^\top. \end{aligned}$$

Here w is the approximate clip and n lies in the normal cone of the spectral ball at w for any B. The error of this resolvent step is

$$\varepsilon = \tfrac{1}{\gamma}(Z - w) - n = \underbrace{\tfrac{1}{\gamma}(I - QQ^\top)Z\hat{V}\hat{V}^\top}_{\varepsilon_{\text{cross}}} + \underbrace{\tfrac{1}{\gamma}(1 - c)T}_{\varepsilon_{\text{tail}}},$$

where $\varepsilon_{\text{cross}}$ is reduced by sweeping the subspace, $B \leftarrow \hat{V}$, and $\varepsilon_{\text{tail}}$ only by increasing k. We write $(w^t, n^t, \varepsilon^t)$ for the output at iterate Z^t.

Fixed-rank DRS DRS with a fixed $k \in \{16, 32, 64, 128\}$ replaces W in (DRS) by $\hat{W}^t = Z^t - \gamma n^t$ and sweeps after every step. Exact DRS uses a full SVD, so that $\varepsilon^t = 0$. A fixed rank must be chosen large enough in advance, since otherwise $\varepsilon_{\text{tail}}$ stalls the iteration.

iDRS iDRS (Pethick et al., 2025a, Alg. II) starts at $k = 4$ and, at iterate Z^t, forms

$$\tilde{Z}^t = w^t - \gamma n^t, \qquad Y^t = \tfrac{1}{\gamma}\big(\tilde{Z}^t - \operatorname{prox}_{\gamma f}(\tilde{Z}^t)\big), \qquad v^t = n^t + Y^t, \qquad \bar{Z}^t = w^t - \gamma Y^t.$$

If the error is small relative to the residual, $\|\varepsilon^t\|_F \le 0.9 \max\{\|v^t\|_F, \|Z^t - \bar{Z}^t\|_F/\gamma\}$, it takes the projective step

$$Z^{t+1} = Z^t - \tfrac{\langle v^t, Z^t - \bar{Z}^t\rangle}{\|v^t\|_F^2}\, v^t.$$

Otherwise it recomputes the spectral clip at Z^t, after doubling k if $\|\varepsilon^t_{\text{tail}}\|_F \ge \|\varepsilon^t_{\text{cross}}\|_F$ and after sweeping the subspace otherwise.

Stopping and timing All methods stop once the duality gap of (17) at $(\mathsf{RowClip}_r(w^t), n^t)$ falls below

$$10^{-6} \min\{s\|G\|_{\text{nuc}},\, r\|G\|_{2,1}\}.$$

Timings are single-threaded NumPy on an Apple-silicon laptop.

Results A fixed rank that is too small stalls at a truncation bias, and the rank needed ranges from 16 to 128 across instances (Table 3). Over all 18 instances, iDRS takes 65 s, against 186 s for exact DRS and 120 s for $k = 128$, the smallest fixed rank that certifies every instance. Choosing the best rank per instance in hindsight would take 54 s. The table also shows that the cheap composed candidates of (16) lose 20–39% of the optimal value on these instances.

Table 3: Joint LMO with inexact spectral steps (1024×256). Median solve time in seconds over three seeds with iDRS's final rank in parentheses. "Best-of-three gap" shows the relative value lost by the best composed candidate U_1, U_2, U_3 in (16). "stall (g)" when DRS fails to certify on every seed, with largest final relative gap g.

			Adaptive / exact		Fixed-rank DRS			
ν	s/r	Best-of-three gap	iDRS	Exact DRS	$k = 16$	$k = 32$	$k = 64$	$k = 128$
1	6	34%	**4.33** (128)	6.33	stall (7e-2)	stall (3e-2)	stall (9e-3)	4.03
1	8	29%	**1.46** (64)	3.76	stall (4e-2)	stall (1e-2)	1.20	2.40
1	12	20%	**0.37** (16)	1.87	0.25	0.37	0.66	1.30
1.5	6	39%	**10.82** (64)	30.79	stall (2e-2)	stall (4e-3)	9.83	19.97
1.5	8	37%	**3.48** (32)	13.82	stall (5e-3)	2.38	4.08	8.75
1.5	12	30%	**1.03** (16)	5.19	0.69	0.94	1.56	3.31

C.4 Spectral clipping via matrix products

For a compact SVD $X = U \operatorname{diag}(\sigma_j) V^\top$, the LMO is $\mathsf{Spectral}_s(X) = sUV^\top$, which Newton–Schulz iterations (Higham, 2008) and Polar Express (Amsel et al., 2025) approximate with matrix products alone, as in Muon (Jordan et al., 2024). Cesista (2025) first expressed $\mathsf{SpectralClip}_s$ through such calls, and Su (2025a,b) derived cheaper expressions. We use the form from

Su (2025b) with two independent calls. For $X \in \mathbb{R}^{m\times d}$ with $m \geq d$, define $P := \frac{1}{2}\big(I + s^{-1}\mathsf{Spectral}_s(X^\top X - s^2 I)\big)$. Then

$$\mathsf{SpectralClip}_s(X) = X - \big(X - \mathsf{Spectral}_s(X)\big)P. \tag{52}$$

In each singular direction, an error in P is weighted by $|\sigma_j - s|$, so errors near the clipping threshold have little effect on the output. A slight excess over s can be removed by rescaling with a certified upper bound on the spectral norm.

D Experimental details

D.1 Robustness-accuracy pareto frontier

Table 4 summarizes the settings of the robustness experiments of Section 7.1.

Table 4: Hyperparameters of the robustness experiments (Section 7.1). $B = \zeta b_0$ is the certificate, and M the momentum of the gradient.

Hyperparameter	SGD + Frobenius	Adam + entrywise	Muon + spectral	Ours + joint
Model	SwiGLU block, $d = 64$, $m = 256$, unconstrained linear head			
Teachers	spread ($\kappa^\star = 0.43$), concentrated ($\kappa^\star = 0.90$)			
Data	$x \sim \mathcal{N}(0, I_d)$, 10 classes, 512 fresh samples per step			
Steps T	1000 (3000 for the long runs)			
Step-size schedule	$\eta_t = \eta_0\big(0.1 + 0.45(1 + \cos(\pi t/T))\big)$			
Head	Adam, step size 10^{-2}, initialized with $\mathcal{N}(0, 1/d)$ entries			
Block initialization	$0.02 \cdot \mathcal{N}(0, 1)$ entries, projected onto the set			
Seeds	0, 1, 2			
Evaluation	4096 fixed held-out samples			
Set of W_g, W_v	Frobenius ball	entrywise box	spectral ball	$\mathcal{K}_{r,s}$
Radius	$c_{\mathrm{F}} = \sqrt{B}$	$c_\infty = \sqrt{B/(d\sqrt{m})}$	$s = \sqrt{B}$	$r = \sqrt{B\kappa^\star}$, $s = \sqrt{B/\kappa^\star}$
Set of W_o	Frobenius ball, radius 1	box, radius $1/\sqrt{dm}$	unit spectral ball	
Direction	$c_{\mathrm{F}} M/\|M\|_F$	$c_\infty \operatorname{adam}(\nabla\mathcal{L})$	$\mathsf{Spectral}_s(M)$	$\mathsf{RowClip}_r(\mathsf{Spectral}_s(M))$
Momentum	$\mu = 0.9$	$\beta_1/\beta_2 = 0.9/0.999$	$\mu = 0.9$	
Step sizes η_0	$\{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$		$\{0.001, 0.003, 0.01, 0.03, 0.1\}$	
Budgets ζ	16 or 64 to 4096; from $\frac{1}{2}$ for seed 0		subsets of $\{\frac{1}{2}, 1, 2, 4, 16, 64\}$	

Task The teacher's head $H^\star \in \mathbb{R}^{10\times d}$ is a Gaussian matrix divided by its spectral norm, so its spectral norm is 1, and labels are the argmax of the teacher's logits. The loss is cross-entropy, and the student's head is trained without a constraint by Adam with step size 10^{-2}.

Teachers Each teacher is generated from a fixed seed as raw matrices (W_o, W_g, W_v) and then normalized. All three are divided by their spectral norms, $\kappa^\star = \max(\|W_g\|_{2\to\infty}, \|W_v\|_{2\to\infty})$ is read off, and W_g, W_v are scaled by $\sqrt{b_0/\kappa^\star}$ with $b_0 = 1/\sqrt{d}$. Both then have spectral norm $\sigma^\star = \sqrt{b_0/\kappa^\star}$, and $\sigma^\star \max(\rho_g^\star, \rho_v^\star) = b_0$, which makes the teacher roughly nonexpansive, and each student is trained at a budget $B = \zeta b_0$. The raw matrices are:

- *Spread* ($d = 64$, $m = 256$): W_g, W_v have i.i.d. $\mathcal{N}(0, 1)$ entries, and $W_o \in \mathbb{R}^{d\times m}$ is a uniformly random matrix with orthonormal rows (from the QR factorization of a Gaussian matrix), so all its singular values are 1.
- *Concentrated* ($d = 64$, $m = 256$): the spread teacher with all but the first 8 rows of W_g and W_v set to zero.

Training Each constrained matrix uses the momentum $M^{t+1} = \mu M^t + (1-\mu)\nabla_W \mathcal{L}$ with $\mu = 0.9$, on 512 fresh samples per step for $T = 1000$ steps, with step size $\eta_t = \eta_0\big(0.1 + 0.45(1 + \cos(\pi t/T))\big)$. Besides $\mathsf{SpectralClip}_s$ and $\mathsf{RowClip}_r$ (Section 6), the projections onto the Frobenius ball and the entrywise box are

$$\mathsf{FrobClip}_{c_\mathrm{F}}(X) = \frac{X}{\max\{1, \|X\|_F/c_\mathrm{F}\}}, \qquad \mathsf{BoxClip}_{c_\infty}(X)_{ij} = \max\{-c_\infty, \min\{X_{ij}, c_\infty\}\}.$$

The updates are

- Muon + spectral ball: $W \leftarrow \mathsf{SpectralClip}_s\big(W - \eta\,\mathsf{Spectral}_s(M)\big)$;
- ours: $W \leftarrow \mathsf{RowClip}_r\big(\mathsf{SpectralClip}_s(W - \eta Z)\big)$ with $Z = \mathsf{RowClip}_r\big(\mathsf{Spectral}_s(M)\big)$ from (16) (Appendix D.1.1 compares it with the direction in Sparrow);
- SGD + Frobenius ball: $W \leftarrow \mathsf{FrobClip}_{c_\mathrm{F}}\big(W - \eta\, c_\mathrm{F} M/\|M\|_F\big)$;
- Adam + entrywise box: $W \leftarrow \mathsf{BoxClip}_{c_\infty}\big(W - \eta\, c_\infty\,\mathrm{adam}(\nabla\mathcal{L})\big)$.

Every step is thus relative to the radius of its set. W_o lies in a set of the same type with worst-case spectral norm 1: the unit spectral ball, Frobenius radius 1, or box radius $1/\sqrt{dm}$. Accuracies and radii are measured on 4096 fixed held-out samples.

Radii of the sets Each set's radius makes its worst-case $\rho\sigma$ equal to B:

- Frobenius ball: $c_\mathrm{F} = \sqrt{B}$;
- entrywise box: $c_\infty = \sqrt{B/(d\sqrt{m})}$, since $\sup \rho\sigma = c_\infty^2 d\sqrt{m}$;
- spectral ball: $s = \sqrt{B}$;
- joint set: $(r, s) = (\sqrt{B\kappa}, \sqrt{B/\kappa})$.

Certified radius The certificate of the trained weights is

$$L = \|H\|_{\mathrm{op}} \tfrac{\sqrt{d}}{2} \|W_o\|_{\mathrm{op}} (L_{\mathrm{SiLU}} \rho_v \sigma_g + \rho_g \sigma_v),$$

from (11) with $\mathbb{N}_{\mathrm{RMS}}$ at $\tau = 1$, for which $\sqrt{d}\,\omega = \sqrt{d}/2$. A correctly classified point is certified at radius margin$/(\sqrt{2}L)$ (Tsuzuku et al., 2018), where the margin is the gap between the largest and second-largest logit: each logit difference $f_y - f_j = (e_y - e_j)^\top f$ is $\sqrt{2}L$-Lipschitz, so a perturbation of smaller norm cannot change the prediction. A misclassified point has radius 0, and the median is taken over all 4096 held-out samples, so it is 0 below 50% accuracy.

Runs Table 4 lists the step sizes, budgets and seeds. The budgets of each method are chosen to cover its accuracy range. The 3000-step runs use seed 0, with peak step sizes $\eta_0 \in \{0.01, 0.03\}$ for Muon and ours. The split sweep of Figure 2c uses seed 0, budgets $\zeta \in \{\frac{1}{2}, 1, 2, 4, 16\}$ and peak step sizes $\eta_0 \in \{0.01, 0.03\}$. Table 5 averages the per-seed frontiers; Figure 2a,b pools the runs of all seeds.

Why the spectral constraint suffers Since W_g is tall, spreading spectral weight evenly over all m units leaves each row a norm of only $s\sqrt{d/m}$ (Remark 5.4), and since $\|a \odot b\|_2 \le \|a\|_\infty \|b\|_2$ for the hidden activation $a \odot b$ with $a = \mathrm{SiLU}(W_g u)$ and $b = W_v u$, the activations are shrunk by the same factor. In order to concentrate the activations, the spectrally constrained weight matrix would therefore need to concentrate the mass on a few singular values. On the spread teacher at $\zeta = \frac{1}{2}$, the spectral ball indeed drives the row-to-spectral ratio of W_g to 0.97. The spectral ball therefore cannot use added width, which we test next.

Table 5: Robustness: maximal clean accuracy and the frontier $R_M(a)$ ($\times 10^{-3}$) at the accuracy used in Section 7.1. Mean $\pm$ standard deviation over three seeds (the 3000-step runs use one seed). "–": accuracy not reached.

	Spread teacher ($\kappa^\star = 0.43$)				Concentrated ($\kappa^\star = 0.90$)	
	1000 steps		3000 steps		1000 steps	
Method	max. accuracy	$R(0.70)$	max. accuracy	$R(0.70)$	max. accuracy	$R(0.85)$
SGD + Frobenius	0.688 ± 0.006	–	0.769	2.76	0.955 ± 0.005	3.01 ± 0.10
Adam + entrywise	0.711 ± 0.007	2.52 ± 0.09	0.780	2.98	0.954 ± 0.003	3.26 ± 0.04
Muon + spectral	0.771 ± 0.003	3.97 ± 0.23	0.835	4.10	0.962 ± 0.003	5.34 ± 0.22
Ours + row/spectral	$\mathbf{0.793 \pm 0.005}$	$\mathbf{4.75 \pm 0.18}$	**0.859**	**5.04**	0.964 ± 0.002	5.25 ± 0.08

Varying width To test the failure mode of a spectral constraint, we consider Gaussian teacher with a large $m^\star = 16d$, normalized as above ($\kappa^\star = 0.25$). Students have width $m \in \{1, 2, 4, 8, 16\} \cdot d$ and use the spectral ball or the joint set with $\kappa = \sqrt{d/m}$, so the two coincide at $m = d$. Each student runs budgets $\zeta \in \{1, 4, 16\}$, peak step sizes $\eta_0 \in \{0.01, 0.03, 0.1\}$ and seeds $0, 1, 2$. Figure 7 shows, at each ζ, the best clean accuracy over η_0 per seed, and the frontier $R(0.70)$ over all ζ and η_0, both averaged over seeds. At $\zeta = 1$, the accuracy of the spectral ball stays flat as the student width grows from $4d$ to $16d$ (0.561 to 0.553), whereas the joint set with $\kappa = \sqrt{d/m}$ improves from 0.680 to 0.729. At $\zeta = 16$, where the budget no longer binds, the plateau disappears.

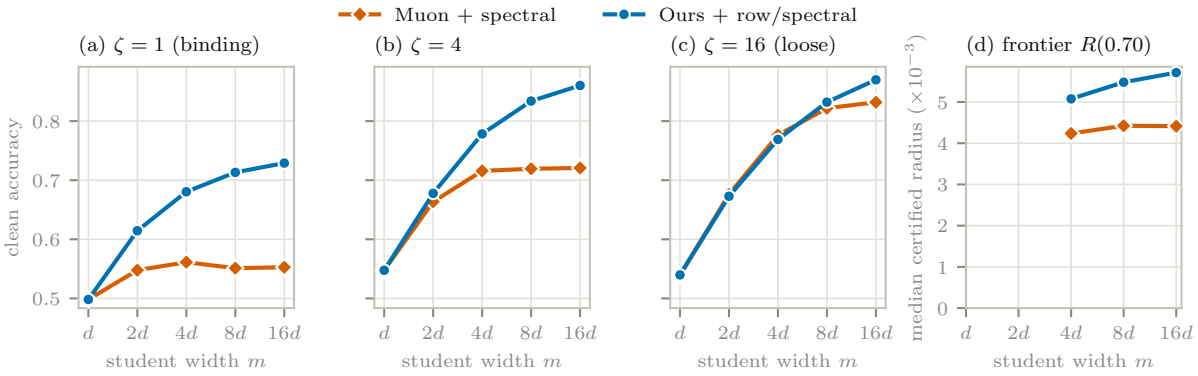

Figure 7: Width sweep at a fixed certificate (teacher $m^\star = 16d$). (a–c) We observe that the spectral constrained method cannot make use of width when the constraint radius is aggressive. (d) The joint constraint certifies a larger radius across widths (at $m \leq 2d$ no run reaches 70%).

D.1.1 Sparrow as the direction of the projected step

The robustness experiments use the direction $U_1 = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s)(M)$ of (16). In contrast, Sparrow uses $U_3 = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s \circ \mathsf{Row}_r)(M)$. Figure 8 compares the two methods using the same budgets, step sizes and seeds as in Section 7.1. For consistency, all directions use knowledge of the teacher's split, $r = \sqrt{B\kappa^\star}$ and $s = \sqrt{B/\kappa^\star}$, instead of the default $r = s\sqrt{d/m}$ in Sparrow.

On the spread teacher the two directions performs practically identically, beating Muon. On the concentrated teacher, however, U_3 performs significantly worse than U_1 and Muon. One could potentially use this as a practical diagnostic. If Sparrow performs worse than Muon in practice, the target might be concentrated on few hidden units. Since a model of adequate size should use its hidden units densely, this in turn suggests that the model is wider than necessary.

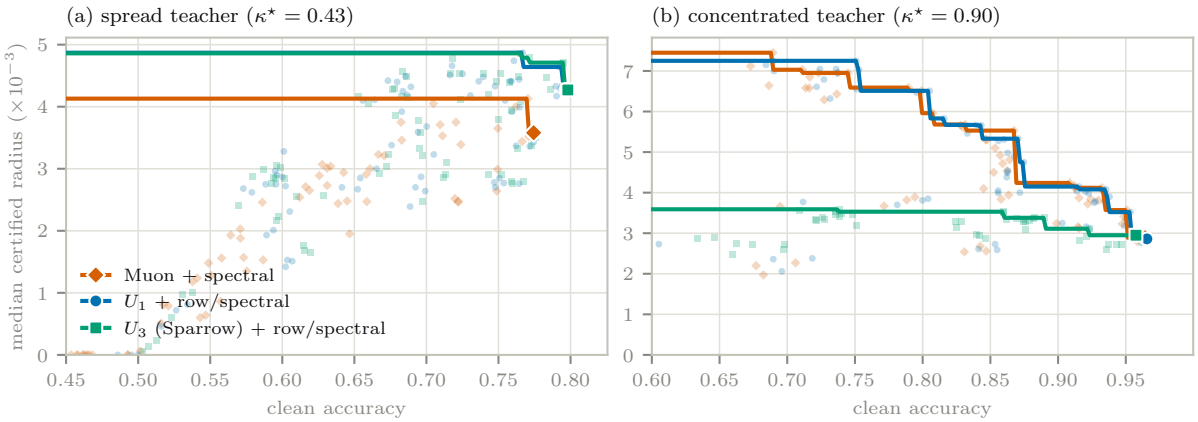

Figure 8: Robustness frontiers $R_M(a)$ of the joint set with the direction U_1, Sparrow's U_3, and Muon. The two directions coincide on the spread teacher, while U_3 falls behind on the concentrated one.

D.2 Sparrow as a drop-in replacement for Muon

Table 6: Hyperparameters of the language-model experiments (Section 7.2). The methods differ only in the direction of the listed matrices. "Sparrow (gates only)" in Figure 4 uses Muon on W_g, W_v and Sparrow on W_a. Step sizes list the values tried for each method (best bolded).

Hyperparameter	Muon	Aurora	Sparrow
Dataset	Tiny Shakespeare, characters, first 90% for training		
Layers / heads / width d	4 / 4 / 128		
Head width	32 (tall gate: 128)		
MLP / normalization	SwiGLU with $m = 512$ / RMSNorm		
Context / batch size	64 / 12		
Steps	1307 (one epoch)		
Step-size schedule	100 warmup steps, then cosine decay to 10%		
Gradient clipping	global norm 1		
Embeddings and head	AdamW, step size 10^{-3}, $\beta_2 = 0.99$, weight decay 0.1		
Initialization	std 0.02, residual projections $0.02/\sqrt{8}$		
Hidden matrices	$W \leftarrow W - \eta\mathsf{lmo}(M)$, no projection, no weight decay		
Momentum	0.9, without Nesterov		
$\mathsf{Spectral}_s$	exact, via SVD (Polar Express in Figure 12)		
Seeds	10		
Evaluation	every 100 steps, 20 batches of the last 10%		
Step size η (Figure 3)	$\{0.005, \mathbf{0.01}, 0.02\}$	$\{0.01, \mathbf{0.02}, 0.04\}$	$\{0.01, \mathbf{0.02}\}$
Step size η (Figure 4)	0.01	0.01	0.01
lmo of W_g, W_v	$\mathsf{Spectral}_1$	Aurora / Sparrow, $s = 1$, $r = s\sqrt{d/m} = \frac{1}{2}$	
lmo of gate W_a	$\mathsf{Spectral}_1$	Aurora if tall	Row_r, $r = s\min\{1, \sqrt{d/(Hd_v)}\} \in \{1, \frac{1}{2}\}$
lmo of other matrices	$\mathsf{Spectral}_1$	Aurora if tall, else $\mathsf{Spectral}_1$	$\mathsf{Spectral}_1$
Weight normalization (hatched)	MuonH	–	SparrowH; without normalized direction
initialization / radius		std $1/\sqrt{n}$ / $R = \|W_0\|_F$	
step size η	$\{0.005, \mathbf{0.01}, 0.02, 0.04, 0.08\}$	–	0.01

Table 6 lists the model and the optimizers. The training split, the first 90% of Tiny Shakespeare, is cut into non-overlapping windows of 64 characters, which are shuffled and each used once, so a run of 1307 steps sees every training character once.

Aurora Aurora is Algorithm 3 of Dewulf et al. (2026) with their default of two iterations and damping $\beta = \frac{1}{2}$: from $X_0 = M/\|M\|_F$ and $D_0 = I$, it computes, for $k = 0, 1$,

$$D_{k+1} = D_k^{\beta}\,\mathrm{diag}\left(\|(X_k)_{i,:}\|_2\right)^{1-\beta}, \qquad X_{k+1} = \mathsf{Spectral}_1\left(D_{k+1}^{-1}X_k\right),$$

and returns sX_2.

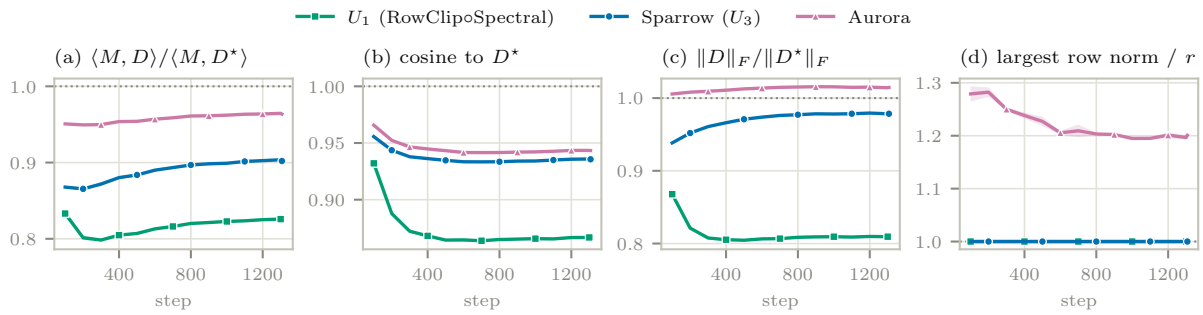

Figure 9: Approximate oracles against the joint LMO $D^\star$ of $\mathcal{K}_{r,s}$ on the momentum of the SwiGLU gate and value matrices along Muon's training runs (mean over the eight matrices, mean $\pm$ one standard error over three seeds). U_1 and Sparrow coincide in (d).

Speed and cost Each method is compared with Muon on the same seed which fixes both the initialization and the data order. We measure speed by the first step at which a run reaches Muon's final validation loss on the same seed. Muon's own validation loss is not monotone, in fact on four seeds it reaches its final value earlier, which makes its mean 1302 rather than 1307. Computing a single direction for a 512×128 matrix takes 2.6 ms for Muon and Sparrow and 5.2 ms for Aurora (mean of 200 compiled calls).

Step-size sweep Figure 10 repeats the comparison over $\eta \in \{0.0025, 0.005, 0.01, 0.02, 0.04\}$ for Muon, U_1 of (16), Aurora and Sparrow on W_g, W_v, with seeds $0, 1, 2$.

Comparison with the exact LMO Along Muon's runs ($\eta = 0.01$, seeds $0, 1, 2$), we map the momentum M of each W_g, W_v by U_1, Sparrow (U_3) and Aurora, and compare each direction D with the joint LMO $D^\star$ of $\mathcal{K}_{r,s}$, computed by DRS to a relative duality gap below $1.1 \cdot 10^{-3}$ (Figure 9). Sparrow attains 87–90% of the optimal value $\langle M, D^\star\rangle$, against 80–83% for U_1. Aurora attains 95–96%, but leaves the set, with a largest row norm 20–28% above r. Importantly, the directions of Sparrow and Aurora are close (cosine to $D^\star$ of 0.93 and 0.94 after the first few hundred steps), so most of Sparrow approximation error is due to length (2–6% below that of $D^\star$), which a larger step size compensates for (Figure 10).

Batch size To vary the batch size without repeating data, Figure 11 uses the Complete Works of Shakespeare (75,365 training windows) and trains for 750 steps at batch $12, 24, 48$ and 96, scoring every run on the same 20×96 validation windows. Each method uses its own best step size per batch size from $\eta \in \{0.01, 0.02\}$ on seeds $0, 1, 2$, with $\eta = 0.04$ added at batch 96, where it loses for all three methods, and $\eta = 0.005$ for Muon at batch 12. The best step size is 0.02 everywhere except for Muon at batch 12, where it is 0.01. The improvement of Sparrow over Muon interesting seems to be larger at small batch sizes.

Inexact polar factor All other runs compute $\mathsf{Spectral}_s$ exactly by an SVD. In Figure 12, every polar factor is instead computed by k iterations of Polar Express (Amsel et al., 2025), at $\eta = 0.02$ for all three methods with seeds $0, 1, 2$. We observe that Sparrow is robust to this approximation of $\mathsf{Spectral}_s$.

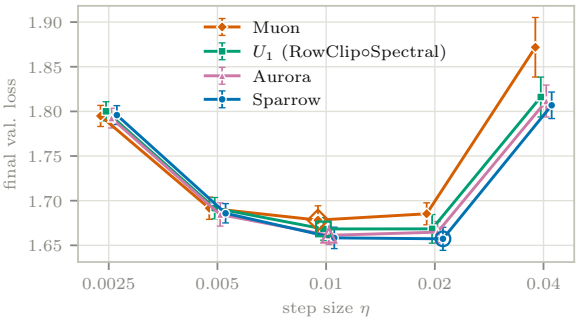

Figure 10: Step-size sweep on Tiny Shakespeare. Final validation loss of Muon, $U_1 = (\mathsf{RowClip}_r \circ \mathsf{Spectral}_s)(M)$, Aurora and Sparrow, which differ only on W_g and W_v. Mean $\pm$ one standard error over three seeds with the best step size of each method being circled. Sparrow seems to allow for slightly larger stepsize.

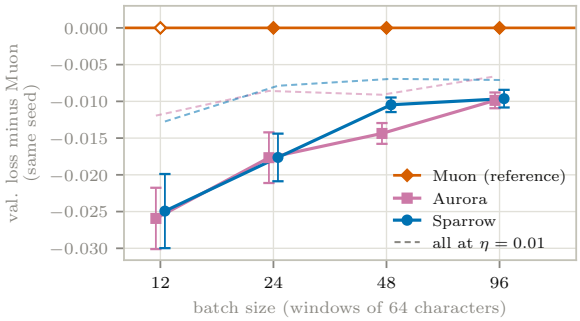

Figure 11: Batch-size sweep on the Complete Works of Shakespeare, 750 steps without repeated data. Each method uses its own best step size per batch size (filled markers: $\eta = 0.02$, hollow: $\eta = 0.01$). Mean $\pm$ one standard error over three seeds. Sparrow is beneficial across batch sizes and particularly at small batch sizes, on par with Aurora at half the cost of its LMO.

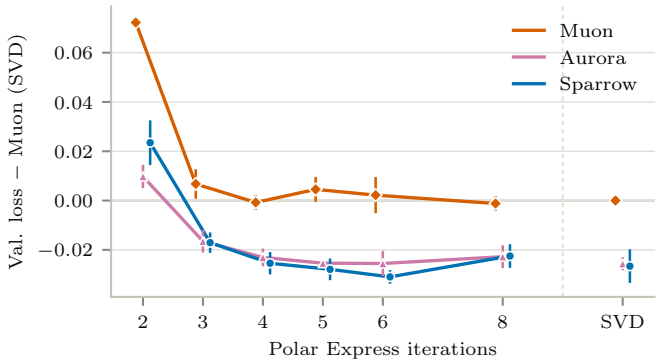

Figure 12: Inexact polar factor on Tiny Shakespeare at $\eta = 0.02$. Final validation loss minus that of Muon with the exact SVD on the same seed, when every polar factor is computed by k iterations of Polar Express (Amsel et al., 2025) and the rightmost points use the SVD. Mean $\pm$ one standard error of the paired difference over three seeds.

Gated attention The gate multiplies the attention output before W_O channelwise by $\sigma(\mathbb{N}(z)W_a^\top)$, with $W_a \in \mathbb{R}^{Hd_v \times d}$, no bias, $c = 1$ and $\mathbb{N}$ the RMSNorm input of the attention block. W_a is initialized from its own random key, so the other weights match the ungated model. Note that the head width 128 of the tall gate also makes W_Q, W_K, W_V tall. Table 6 gives the directions and radii of each method.

Weight normalization For the hatched bars of Figure 4, every hidden matrix $W \in \mathbb{R}^{m \times n}$ is initialized with entries of standard deviation $1/\sqrt{n}$ and kept at $R = \|W_0\|_F$ by N of (50) (Wen et al., 2026). With D the direction of Muon or Sparrow, MuonH and SparrowH take $W \leftarrow \mathsf{N}(W - \eta\,\mathsf{N}(D))$, and Sparrow without normalized direction takes $W \leftarrow \mathsf{N}\big(W - \eta \frac{R}{s\sqrt{\min\{m,n\}}} D\big)$. Since the initialization differs from the other runs, the effect of Sparrow under weight normalization should be compared against MuonH.

D.3 Looped models

The looped model of Section 7.2 is the tall-gate model of Table 6. The first layer maps the embedding to the anchor u, the second is iterated N times, and the third reads out z^N:

$$\begin{aligned}
\text{prelude:} \quad & u = (\mathbb{I} + \mathbb{M}_1) \circ (\mathbb{I} + \mathbb{A}_1)(E_x + P), \\
\text{loop:} \quad & z^{k+1} = \tfrac{1}{2}u + \tfrac{1}{2}\mathbb{T}(z^k), \qquad z^0 = u, \quad k < N, \\
\text{core:} \quad & \mathbb{T} = (\mathbb{I} + \mathbb{M}_2) \circ (\mathbb{I} + \mathbb{A}_2), \\
\text{coda:} \quad & \text{logits} = \mathbb{N}\big((\mathbb{I} + \mathbb{M}_3) \circ (\mathbb{I} + \mathbb{A}_3)(z^N)\big) W_{\text{head}}^\top,
\end{aligned}$$

where $E_x = (E_{x_1}, \ldots, E_{x_T})$ are the rows of the learned token embedding E for the input characters and P the learned position embeddings, $\mathbb{A}_\ell$ is the gated attention and $\mathbb{M}_\ell$ the SwiGLU block (10) of layer ℓ, both with the RMSNorm $\mathbb{N}$ on their input. Truncated backpropagation detaches z^{N-H} (the gradient still flows to the first layer through the anchor u). Sparrow is applied to all matrices, as in Sparrow (all) of Figure 4. Each method uses its best step size of $\{0.01, 0.02\}$ per configuration across three seeds: 0.01 for Muon (except 0.02 at $H = 8$) and 0.02 for Sparrow (except 0.01 at $N = 1$).