Solve the Loop: Attractor Models for Language and Reasoning

Jacob Fein-Ashley [1] **Paria Rashidinejad** [1]

Abstract

Looped Transformers offer a promising alternative to purely feed-forward computation by iteratively refining latent representations, improving language modeling and reasoning. Yet recurrent architectures remain unstable to train, costly to optimize and deploy, and constrained to small, fixed recurrence depths. We introduce *Attractor Models*, in which a backbone module first proposes output embeddings, then an attractor module refines them by solving for the fixed point, with gradients obtained through implicit differentiation. Thus, training memory remains constant in effective depth, and iterations are chosen adaptively by convergence. Empirically, Attractor Models outperform existing models across two regimes, large-scale language-model pretraining and reasoning with tiny models. In language modeling, Attractor Models deliver a *Pareto improvement* over standard Transformers and stable looped models across sizes, improving perplexity by up to 46.6% and downstream accuracy by up to 19.7% while reducing training cost. Notably, a 770M Attractor Model outperforms a 1.3B Transformer trained on twice as many tokens. On challenging reasoning tasks, we show that our model with only 27M parameters and approximately 1000 examples achieves 91.4% accuracy on Sudoku-Extreme and 93.1% on Maze-Hard, scaling favorably where frontier models like Claude and GPT o3, fail completely, and specialized recursive reasoners collapse at larger sizes. Lastly, we show that Attractor Models exhibit a novel phenomenon, which we call *equilibrium internalization*: fixed-point training places the model's initial output embedding near equilibrium, allowing the solver to be removed at inference time with little degradation. Together, these results suggest that Attractor Models make iterative refinement scalable by turning recurrence into a computation the model can learn to internalize.

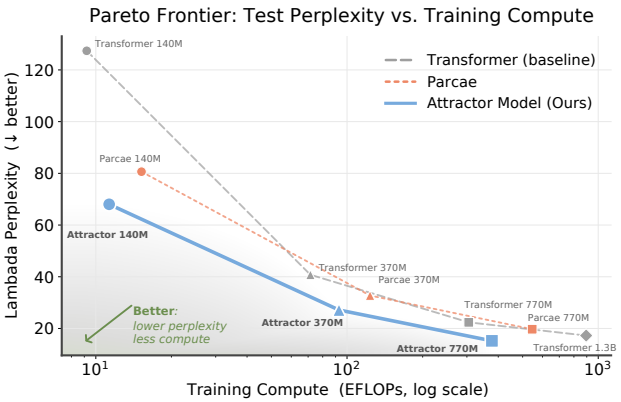

Figure 1. Pareto frontier of Lambada perplexity vs training FLOPs. Attractor Models achieve strong performance on large-scale language modeling with less compute.

1. Introduction

The modern language-modeling era has been dominated by Transformers (Vaswani et al., 2017), which produce each token through a fixed feed-forward computation. This recipe has been extraordinarily successful (OpenAI, 2023; Gemini Team, 2023; Grattafiori et al., 2024; Anthropic, 2024; Guo et al., 2025), but it leaves a basic question unresolved: should each token be the product of a single pass of computation, or should a model be able to refine its latent prediction before committing to an output? A growing body of work suggests that such refinement can be powerful. Chain-of-thought reasoning (Wei et al., 2023) can be viewed as one form of such refinement, where a model writes intermediate tokens, feeds them back into its context, and uses them to shape later predictions. Yet this routes computation through the discrete token channel and forces "thinking" to be written down, even when the effect might be to merely refine internal representations.

This limitation has inspired several lines of work on latent (or implicit) thinking and a re-emergence of architectural recurrence, which move thinking away from purely token-level generation. These include universal Transformers (Dehghani et al., 2019), looped Transformers (Giannou et al., 2023a; Yang et al., 2024a), recurrent-depth Transformers

[1]University of Southern California. Correspondence to: Jacob Fein-Ashley <feinashl@usc.edu>.

Published as a paper at the 1st *FoGen workshop, ICML 2026*, Seoul, South Korea, 2026. Copyright 2026 by the author(s).

(Kohli et al., 2026), looped language models (Zhu et al., 2025c), latent reasoning methods (Geiping et al., 2025a; Saunshi et al., 2025a), and continuous chain-of-thought approaches (Hao et al., 2025; Mohtashami et al., 2023; Zhu et al., 2026). Looped architectures can, in principle, express iterative or algorithmic procedures (Yang et al., 2024b; Giannou et al., 2023a), emulate additional depth through weight sharing (Dehghani et al., 2019; Zhu et al., 2025c), reduce the context-length costs of token-level reasoning, and improve downstream generalization (Labovich, 2026b; Fan et al., 2025). Empirically, recent looped language models offer gains in language modeling and reasoning (Zhu et al., 2025c; Geiping et al., 2025a), and tiny recursive models (Wang et al., 2025; Jolicoeur-Martineau, 2025) have shown that recurrence can be useful in hard reasoning tasks in small-data regimes.

The challenge is that recurrence has proven difficult to use as a stable architectural building block. Recurrence is often accompanied by unstable training, large memory requirements that grow linearly with the number of recurrent steps, and significant, sequential compute (Geiping et al., 2025a; Zhu et al., 2025c; Prairie et al., 2026). Training recurrent networks typically requires backpropagation through time (or, depth) and carefully designed stabilization techniques; even then, latent-thinking models remain fragile and difficult to optimize (Wei et al., 2025; Özeren & Aßenmacher, 2025; Deng et al., 2026; 2025; Rizvi-Martel et al., 2026). For training, looped language models tend to require substantially more compute than comparable feed-forward models and can become memory-limited at larger recurrence depths. For instance, Geiping et al. (2025a) reports that training a recurrent model can consume raw FLOPs comparable to those of a feed-forward model ten times larger. At the opposite end of the spectrum, specialized tiny recursive reasoners exhibit a troubling "less is more" behavior and respond negatively to scaling: increasing model size can degrade or even collapse performance (Jolicoeur-Martineau, 2025).

1.1. Contributions

In this work, we design a general-purpose architecture for iterative refinement that is (i) stable to train, (ii) uses constant-memory in the number of refinement steps, (iii) is substantially cheaper to train than explicit unrolling, (iv) is efficient during inference, and (v) achieves a strong performance across both large-scale language modeling and hard reasoning with tiny models.

Refine outputs by solving the loop with Attractor Models. We introduce *Attractor Models*, a new family of architectures that treat latent refinement as a fixed-point problem in the output embedding space. The model first proposes an initial guess embedding using a non-recurrent backbone module (implemented as a Transformer in ours). A separate, typically smaller, recurrent network then refines this guess. Recent mechanistic analyses of looped language models demonstrate that, for the vast majority of tokens, the recurrent trajectory converges to a fixed point (Blayney et al., 2026). We build directly on this observation and instead of unrolling the loop for a predefined number of steps, we solve for the state to which the loop converges, inspired by Deep Equilibrium Models (DEQ; (Bai et al., 2019)). The name Attractor Model comes from dynamical systems, where an attractor is a set of states toward which a system evolves. In a sense, Attractor Models can be viewed as *thinking* before producing each token: the backbone proposes an initial latent prediction, the attractor module refines it to equilibrium, then decodes it into the output distribution.

Attractor Models offer stable, constant-memory, efficient training, and adaptive refinement. Unlike looped LMs, which finitely unroll the recurrent block, Attractor Models solve an equilibrium by treating the prediction target as a fixed-point computation. The number of refinement steps is therefore chosen adaptively according to convergence during both training and inference. We show that the memory cost during training remains constant in the number of iterations; whereas standard looped language models have a linear scaling increase with the number of loops. Our experiments demonstrate that the two-stage structure of Attractor Models, in which the backbone proposes and the attractor refines, enables stable, efficient training and strong performance.

Novel phenomenon: Equilibrium internalization. We observe that despite the fact that Attractor models are trained only with the next-token prediction loss, they learn to make the solver unnecessary. During training, the backbone's initial prediction moves progressively closer to the fixed point, so fewer refinement steps are needed to reach approximate equilibrium (c.f. Figures 6 and 5). We call this phenomenon *equilibrium internalization*: the model appears to self-distill the iterative refinement process into its own initial output embedding, through a form of automatic curriculum. In this sense, recurrence acts as a moving training target, teaching the backbone where its computation should converge.

Strong performance in large-scale language modeling and hard reasoning with tiny models. Our experiments show that Attractor Models scale across two regimes. In large-scale language modeling, Attractor Models consistently outperform standard Transformers and stable looped language models across small (140M), medium (370M), and large (770M) sizes, delivering a Pareto improvement. We show that our models improve validation perplexity, out-of-distribution perplexity on Lambada (Paperno et al., 2016), and downstream benchmark accuracy while using substantially less training compute than comparable looped baselines. Notably, a 770M-parameter Attractor Model out-

performs a 1.3B-parameter Transformer trained on twice as many tokens. Compared to looped LM Parcae (Prairie et al., 2026), our models use up to 31% less training compute, while avoiding the memory growth associated with explicit unrolling. In hard reasoning tasks with tiny models, with only 27M parameters and approximately 1000 training examples, Attractor Models achieve 91.4% accuracy on Sudoku-Extreme and 93.1% on Maze-Hard. In this regime, standard Transformers as well as proprietary frontier models such as DeepSeek R1, Claude, and o3-mini fail completely at 0%, while specialized recursive architectures underperform our model and collapse when scaled. Attractor Models, in contrast, improve with scale.

2. Background: Looped Architectures

We begin with background on looped architectures. Let $x = (x_1, \ldots, x_n) \in \mathcal{V}^n$ be an input sequence over vocabulary $\mathcal{V}$, and let d denote the model width. Looped models can be written as a composition of three units: a prelude unit $\tilde{x} = \mathcal{P}(x) \in \mathbb{R}^{n \times d}$, which produces an input representation $\tilde{x} \in \mathbb{R}^{n \times d}$; a weight-tied recurrent unit $h_{t+1} = \mathcal{R}(h_t, \tilde{x})$, which is applied repeatedly to a latent state $h_t \in \mathbb{R}^{n \times d}$ for T steps; and a coda unit, which maps the final latent state to output probabilities $p = \mathcal{C}(h_T) \in \Delta(\mathcal{V})^n$. Importantly, looped architectures commonly initialize the latent state at an *uninformative value*, such as $h_0 = 0$ or Gaussian noise $h_0 \sim \mathcal{N}(0, \sigma^2 I)$ (Geiping et al., 2025a; Prairie et al., 2026; Bai et al., 2019). Furthermore, the recurrent step may use the input representation only at the first step (Zhu et al., 2025c). or at every recurrent step (Geiping et al., 2025a; Prairie et al., 2026). Such injection may be through addition or concatenation with the recurrent state.

Models such as Parcae (Prairie et al., 2026), Huggin (Geiping et al., 2025a), and Ouro (Zhu et al., 2025c) differ mainly in how they train, stop, or scale this looped architecture. In particular, the recurrence depth T is a central design choice in these models. It may be fixed (Jolicoeur-Martineau, 2025), sampled during training (Geiping et al., 2025a; McLeish et al.; Prairie et al., 2026), or determined by an auxiliary halting mechanism (Bae et al., 2025b; Zhu et al., 2025c). Training then minimizes an objective averaged over both the data distribution and the chosen recurrence-depth mechanism, typically by backpropagation through depth. Consequently, both training cost and gradient memory are tied to the number of recurrent steps. Furthermore, changing T at inference introduces a train–test mismatch, since the model is evaluated under a different computation graph than the one used during training, leading to degraded performance.

3. Solve the Loop with Attractor Models

As discussed in the previous section, standard looped language models (Zhu et al., 2025c; Prairie et al., 2026) use weight sharing to recurrently refine a hidden state that is initialized from an uninformative value and based on input embeddings. Predictions are then read out after a finite number of loops (Prairie et al., 2026), or once an auxiliary halting head becomes confident (Graves, 2017; Zhu et al., 2025c; Prairie et al., 2026). This design carries three drawbacks: the loop count must be chosen at training time, training memory grows linearly in the number of loops, and accuracy degrades when more loops are run at inference than were seen during training (Zhu et al., 2025c). As a result, recurrence often comes with unstable training, growing memory requirements, and large sequential compute, in some cases approaching the costs of training non-recurrent models ten times larger (Geiping et al., 2025a).

Recent mechanistic analyses of looped language models (Blayney et al., 2026) reveals that for the vast majority of tokens, the recurrent trajectory eventually converges to a fixed point. This suggests that the weight-tied recurrent modules are often approximating an underlying fixed-point computation, doing so through the recursive application of the weight-tied block truncated after T steps. This observation motivates the design of Attractor Models, which we subsequently describe.

3.1. Attractor Model: Backbone and Attractor Modules

Motivated by the fixed-point behavior observed in looped models, we model recurrent refinement as an attractor. Rather than training a model to produce good predictions after a prescribed number of recurrent steps, we define the output as the equilibrium of the refinement process. Attractor Models consist of two modules: the *backbone module* (typically a larger Transformer network) first proposes a meaningful initial output embedding, and the *attractor module* (typically a smaller Transformer-based network) then refines this proposal until convergence. This makes the number of refinement steps a solver choice rather than a fixed architectural choice.

We first start by mapping the inputs x into input embeddings $\tilde{x} = E(x) \in \mathbb{R}^{n \times d}$, where E denotes the tied embedding/unembedding. Then, the input embedding is processed by the backbone and attractor modules as described below.

The backbone module proposes an initial "guess" output embedding. The backbone module $\mathcal{T}_{\theta_b}$ maps the input embeddings to an initial proposal:

$$\tilde{y}_0 = \mathcal{T}_{\theta_b}(\tilde{x}), \quad \text{where} \quad \tilde{x} = E(x). \tag{1}$$

We use $\tilde{y}_0$ as an initialization for the attractor module. Instead of initializing the loop from zero, noise, or an input-

side representation, Attractor Models initialize the recurrent computation from a state that is already a coherent prediction embedding. In practice, $\mathcal{T}_{\theta_b}$ is a relatively high-capacity causal Transformer, so the refinement begins near a meaningful initialization rather than 0. We find that this makes training our method stable compared to DEQ, which experiences a blow-up in the number of iterations used later in training; whereas our method stabilizes later in training (c.f. Figure 6(b)).

The attractor module refines the output embedding. The attractor module is a separate weight-tied refinement network $\mathcal{T}_{\theta_a}$. Starting from the backbone proposal $\tilde{y}_0$, it repeatedly refines the output embedding according to

$$\tilde{y}_{t+1} = \mathcal{T}_{\theta_a}(\tilde{y}_t, \tilde{y}_0), \quad \text{where} \quad \tilde{y}_0 = \mathcal{T}_{\theta_b}(\tilde{x}). \tag{2}$$

Here, we persistently inject the initial guess $\tilde{y}_0$ at every refinement step. This persistent injection keeps the attractor proposal-dependent and prevents it from collapsing to a proposal-independent fixed point. We ablate this conditioning mechanism in Section 4. Importantly, we *warm-start* the attractor module at an informative proposal $\tilde{y}_0$, in contrast to existing work that initialize the recurrent state at uninformative values such as zero or Gaussian noise (Geiping et al., 2025a; Prairie et al., 2026).

Rather than rolling out recurrent steps to reach a fixed point, we directly solve for the convergence:

$$\mathcal{A}_{\theta_a}(\tilde{y}^\star, \tilde{y}_0) \coloneqq \mathcal{T}_{\theta_a}(\tilde{y}^\star, \tilde{y}_0) - \tilde{y}^\star = 0, \tag{3}$$

$$\Rightarrow \quad \tilde{y}^\star = \texttt{RootFind}(\mathcal{A}_{\theta_a}(\cdot, \tilde{y}_0); y_0)\,. \tag{4}$$

In the forward pass, we compute this equilibrium with a root finder initialized at the backbone proposal. In our implementation, the `RootFind` algorithm uses Anderson acceleration, which combines a small window of past iterates and residuals to reach the fixed point faster than plain recursion. The solver exits when $\|\mathcal{A}_{\theta_a}(\tilde{y}_t, \tilde{y}_0)\|_2 / \|\tilde{y}_t\|_2 < \varepsilon$ or after $T_{\max}$ steps. Thus, the computation is controlled by the convergence of the residual rather than by a learned halting head or a preset loop count. In contrast to fixed unrolling, the number of refinement steps can therefore vary at inference time without changing the model. Finally, the equilibrium embedding is decoded with the tied unembedding.

Parameters of the Attractor Models consist of the tied embedding/unembedding matrices, the backbone module, and parameters of the attractor module: $\theta \coloneqq (\theta_a, \theta_b, E)$. Compared to looped models, Attractor Models change both the starting point and the endpoint of recurrence: we initialize the loop from the output guess from the backbone network $\tilde{y}_0$, and the decoded state is the attractor $\tilde{y}^\star$ rather than a finite unroll.

3.2. Training and Inference of Attractor Models

We now describe the training procedure for Attractor Models. We first explain how to differentiate through the fixed-point solver using implicit differentiation, and then show how the model is optimized with the standard cross-entropy language-modeling loss applied to the output $y^\star$.

Backward pass and implicit differentiation. Because Attractor Models define the output embedding as the solution to a fixed-point equation, we differentiate through the equilibrium using the implicit function theorem (Krantz & Parks, 2002b). Let $\mathcal{L}$ denote the training loss and let $v = \partial\mathcal{L}/\partial\tilde{y}^\star$. Applying the implicit function theorem to $\mathcal{A}_{\theta_a}(\tilde{y}^\star, \tilde{y}_0) = 0$ gives

$$\frac{\partial \mathcal{L}}{\partial \theta} = u^\top \frac{\partial \mathcal{T}_{\theta_a}(\tilde{y}^\star, \tilde{y}_0)}{\partial \theta}, \quad u = \left(I - J_{\tilde{y}}^\top\right)^{-1} v, \tag{5}$$

$$\text{where} \quad J_{\tilde{y}} = \left.\frac{\partial \mathcal{T}_{\theta_a}(\tilde{y}, \tilde{y}_0)}{\partial \tilde{y}}\right|_{\tilde{y}=\tilde{y}^\star}. \tag{6}$$

The derivative with respect to $\theta = (\theta_a, \theta_b, E)$ includes the direct dependence on the attractor parameters θ_a, as well as the dependence on the initialization $\tilde{y}_0 = \mathcal{T}_{\theta_b}(E(x))$ through the backbone parameters θ_b and the tied embedding parameters E.

Following prior work on implicit models (Geng et al., 2022; Fung et al., 2021), we use the one-step approximation $u \approx v$. This avoids the extra linear solve for u and reduces the backward pass to one vector–Jacobian product through $\mathcal{A}_{\theta_a}$. Since we do not backpropagate through every solver step, memory in the attractor block does not grow with the number of forward iterations. In Section 4, we show that the Anderson solver yields only marginal quality gains, whereas the one-step approximation enables substantially cheaper training.

Remark. Attractor Models are implicit equilibrium models in the spirit of DEQs (Bai et al., 2019), but the equilibrium plays a different role. Classical DEQs replace the prediction network with a hidden-state equilibrium $z^\star$ (single-layer), decoded with a separate output head. We instead keep a standard causal Transformer backbone and add an equilibrium refinement block on top of its prediction state. The fixed point lives directly in the tied embedding space, so every iterate $\{y_0, y_1, \ldots, y^\star\}$ is already a representation in the output space that can be decoded. This gives three practical differences: (i) the solver is initialized from a semantically meaningful proposal $\tilde{y}_0$ rather than from an uninformative state such as zero (as in DEQ), (ii) inference can stop according to a residual tolerance ε rather than a fixed depth or learned halting head, and (iii) DEQ shows that scaling the number of DEQ blocks can harm the performance of their method, whereas we allow for an arbitrary depth backbone transformer and show that we can use a variable number of solver blocks.

Training objective and inference. We train Attractor Models with the standard next-token prediction cross-entropy objective applied to the fixed-point output $y^\star$. Inference reuses the same equilibrium computation. Given an input sequence x, the backbone first produces the proposal $\tilde{y}_0 = \mathcal{T}_{\theta_b}(E(x))$, the attractor solver computes $\tilde{y}^\star$, and the tied unembedding decodes $\tilde{y}^\star$ into next-token probabilities. Peak memory is bounded by a single forward through the attractor module and standard KV-caching applies in the backbone. In principle, the solver tolerance ε and maximum iteration budget $T_{\max}$ are inference-time hyperparameters: they can be adjusted without retraining the model, turning test-time computation into a budget for approaching the learned attractor. Interestingly, however, we find that trained Attractor Models often require very little test-time refinement, as we describe in the next section.

3.3. Equilibrium Internalization and Stability in Attractor Models

Although Attractor Models define predictions through the equilibrium $\tilde{y}^\star$, we observe a surprising phenomenon: after training, the backbone proposal $\tilde{y}_0$ often lies close to the equilibrium (c.f. Figures 6 and 5). We refer to this phenomenon as *equilibrium internalization*. Intuitively, the attractor module appears to act as a moving teacher for the backbone, resulting in a form of automatic curriculum. Early in training, the proposal $\tilde{y}_0$ may be far from a good prediction, and the solver must perform nontrivial refinement to reach $\tilde{y}^\star$. Since $\tilde{y}_0$ and $\tilde{y}^\star$ live in the same tied output-embedding space, gradients through the equilibrium also train the backbone proposal to move toward the state that the solver would have found. Thus, when the backbone is sufficiently expressive, much of the prediction work can be internalized into $\tilde{y}_0$, leaving the attractor to perform a stable refinement.

Stability. Equilibrium training also biases the recurrent map toward convergent dynamics. The implicit gradient contains the inverse factor $(I - J_{\tilde{y}}^\top)^{-1}$, which becomes ill-conditioned near non-contractive regimes. This creates a barrier against unstable fixed-point dynamics, unlike fixed-loop training, which can learn trajectories that are accurate only at a prescribed step count and fail under extra inference-time loops. We discuss this contrast in detail in Appendix B.3. We refer to Appendix B for theoretical analysis of Attractor Models and comparison with finite-loop models.

4. Experiments

In this section, we evaluate Attractor Models in two regimes. We first study large-scale language modeling across model sizes, comparing against parameter-matched Transformers and looped-LM baselines. We evaluate scaling behavior, downstream accuracy, and training efficiency. We also present results on hard reasoning tasks, where we test whether the same fixed-point refinement mechanism improves small models on problems that require iterative computation.

4.1. Attractor Models Improve Large-Scale Language Modeling

Setup. We follow the `nanochat` (Karpathy, 2025) pretraining recipe used by Parcae (Prairie et al., 2026) for its main Transformer comparison, training on FineWeb-Edu (Penedo et al., 2024). To ensure a fair comparison, all models are matched in parameter count and trained with the same data budget, optimizer, and learning-rate schedule as the Parcae baselines; the only change is the recurrent block. We compare at three scales: 140M, 370M, and 770M parameters. Parcae is a middle-looped language model with prelude, coda, and recurrent blocks. The stability in this model comes from a linear injection that bounds the spectral radius of the recurrence below one. In contrast, our architecture uses standard Transformer blocks followed by a fixed-point iteration block, and lets the solver itself control the effective depth. Detailed hyperparameter settings are in Appendix C.

Parameter Scaling. We demonstrate how large-scale pretraining scales with our model. We evaluate our method on 140M, 370M, and 770M parameters against a parameter-matched Transformer and PARCAE, a looped-language model (Prairie et al., 2026). Our model improves monotonically with scale. Across all three sizes, our method achieves the best validation PPL, Lambada PPL, and CORE accuracy. These results show that our fixed-point refinement scales cleanly with model size, with especially large gains on Lambada, where iterative refinement substantially improves long-context prediction.

Figure 2. **Training-time efficiency.** Despite being a looped architecture, our method uses 25–31% fewer FLOPs than Parcae because the fixed-point solver often converges before $T_{\max}$ and the backward pass uses the one-step implicit-gradient approximation.

Training efficiency: Lower FLOPs. The IFT backward pass keeps training memory constant in the number of solver iterations. The memory of standard looped language models like (Prairie et al., 2026) scales linearly with the number of loops. The total training FLOPs (Figure 2) follow the

Table 1. Parameter scaling results.

Size	Model	Val. PPL ↓	Lambada PPL ↓	Core ↑	Core-Ext. ↑
140M	Transformer	21.48	127.39	13.00 ± 0.15	8.80 ± 0.21
	PARCAE	19.06	80.64	14.04 ± 0.20	9.67 ± 0.28
	Attractor Model	**18.30**	**68.02**	$\mathbf{14.59 \pm 0.11}$	$\mathbf{10.03 \pm 0.05}$
	rel. Δ	↓14.8%	↓46.6%	↑12.2%	↑14.0%
370M	Transformer	15.79	40.77	17.46 ± 0.03	11.71 ± 0.22
	PARCAE	14.49	32.74	20.00 ± 0.06	$\mathbf{12.75 \pm 0.31}$
	Attractor Model	**14.03**	**27.14**	$\mathbf{20.24 \pm 0.09}$	12.64 ± 0.33
	rel. Δ	↓11.1%	↓33.4%	↑15.9%	↑7.9%
770M	Transformer	13.08	22.37	22.42 ± 0.20	14.20 ± 0.63
	PARCAE	12.49	19.71	25.07 ± 0.33	15.19 ± 0.43
	Attractor Model	**12.09**	**15.21**	$\mathbf{26.83 \pm 0.29}$	$\mathbf{15.42 \pm 0.51}$
	rel. Δ	↓7.6%	↓32.0%	↑19.7%	↑8.6%
1.3B	Transformer	11.95	17.26	25.45 ± 0.08	15.90 ± 0.23

same trend: although every step of our recurrent block costs roughly the same as Parcae's, our solver typically converges below ε in well under $T_{\max}$ steps, so the realized depth during training is lower despite identical $T_{\max}$, yielding 25–31% lower training FLOPs across scales.

Training efficiency: O(1) memory An additional advantage of our method is that the training memory required stays constant with the number of iterations (Figure 3). This is because our implicit backward pass does not need to store the intermediate activations from every recurrent step. In contrast, standard looped language models must backpropagate through each unrolled iteration, causing memory usage to grow linearly with the number of loops.

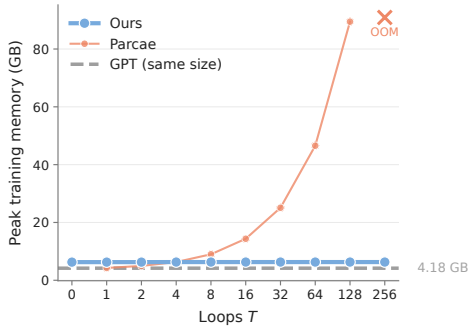

Figure 3. Training memory requirement varying the number of loops/iterations.

4.2. Attractor Models Lead to Significant Gains in Hard Reasoning Tasks with Tiny Models

Beyond language modeling, we evaluate on Sudoku-Extreme and Maze-Hard (Wang et al., 2025; Jolicoeur-Martineau, 2025): two challenging reasoning benchmarks, where standard Transformers and most frontier LLMs fail.

Setup. We train small models with 1000 training examples for each task and require predicting the full output grid in a single direct forward pass (no autoregressive decoding). We follow the TRM training protocol (Jolicoeur-Martineau, 2025), using deep-supervision steps.

Method. TRM carries two latents across deep-supervision steps, a current answer y and a reasoning state z, and applies a tiny two-layer network $T(n+1)$ times per step to update them. We keep this protocol except for the inner update: instead of unrolling $T(n+1)$ applications, we solve directly for the fixed point of the (y, z) update with our solver. The initialization is handled by deep supervision itself: the previous step's (y, z) initializes the solver at the next step, with a learned embedding at step zero, so we do not use a separate backbone.

We present our results in Table 2. The fixed-depth Transformer fails on both tasks. HRM (27M) achieves 55.0% / 74.5%. TRM is the strongest tiny baseline at 7M (74.7% / 85.3%), but (counter to the goal of scaling) collapses to 0% on both tasks when scaled to 27M parameters. Our model scales naturally with parameter count. We attribute the difference to the explicit fixed-point objective, which appears to provide regularization that bare iterative refinement lacks at higher capacity.

For the backward pass we use the phantom-gradient scheme (Geng et al., 2022) instead of the one-step approximation used in the language-modeling experiments. We find this necessary at the parameter counts considered here. The very small networks and training sets ($\sim$1,000 examples) are too sensitive to the one-step approximation, consistent with TRM's report that switching to one-step drops Sudoku-Extreme accuracy from 87.4% to 56.5% (Jolicoeur-Martineau, 2025).

This setup is still an instance of our Attractor Model framework, where an output-space representation is iteratively refined to an equilibrium, decoded, and differentiated through implicitly. The only difference is the initialization mechanism. Rather than a Transformer backbone producing $\tilde{y}_0$, deep supervision supplies the initialization. Specifically, the previous supervision step's (y, z) initializes the solver at the next step, with a learned embedding at step zero.

Table 2. Small-sample training results on Sudoku-Extreme and Maze-Hard.

Method	# Params	Sudoku ↑	Maze ↑
Deepseek R1	671B	0.0%	0.0%
Claude 3.7	?	0.0%	0.0%
O3-mini-high	?	0.0%	0.0%
Transformer	27M	0.0%	0.0%
HRM	27M	55.0%	74.5%
TRM	7M	74.7%	85.3%
TRM	27M	0.0%	0.0%
Ours	7M	54.3%	46.7%
Ours	27M	**91.4%**	**93.1%**

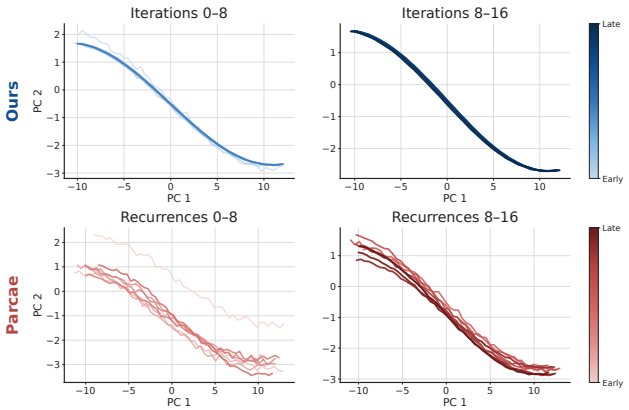

Figure 4. PCA projection of the residual stream at the final sequence position over 16 iterations.

4.3. Equilibrium Internalization and Test-Time Behavior

Fixed-point convergence. Figures 4 and 6 visualize convergence in two complementary ways. First, we project the trajectory of the final-position representation onto its first two principal components over 16 iterations. Our method rapidly contracts to a fixed point: iterations 8–16 collapse onto a single attractor. Parcae also moves toward a stable state, but does so more slowly; its trajectory remains noisier and continues to drift through later recurrences. This suggests that while finite-loop language models can exhibit fixed-point-like behavior, explicitly training the loop as a fixed-point problem produces faster and cleaner convergence. Second, we track the number of solver iterations required during training. The DEQ baseline requires increasingly many iterations as optimization progresses, consistent with the observations in the original work (Bai et al., 2019). In contrast, Attractor Models rapidly converge to the minimum iteration count and remain stable. This behavior provides evidence for equilibrium internalization, where optimization shifts work from the iterative solver into the backbone proposal.

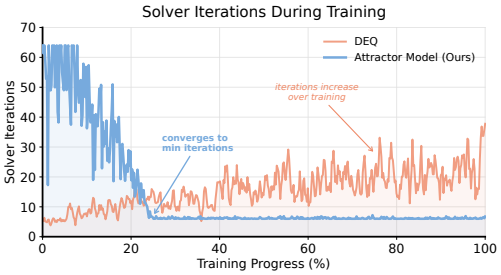

Figure 6. The DEQ baseline requires increasingly many solver iterations during training, whereas ours rapidly settles to the minimum iteration count and remains stable.

Test-time iterations vs. quality. In Figure 5, we report validation perplexity, CORE accuracy, and CORE-Extended accuracy as we vary the number of inference iterations T while holding training fixed. For Parcae, quality improves monotonically from $T = 1$ until it plateaus near $T = 8$, indicating that the model relies on repeated test-time refinement. In contrast, our method reaches peak performance at $T = 1$ at every scale, and even $T = 0$ (decoding the backbone proposal $\tilde{y}_0$ directly through the tied unembedding) is already at or near the converged value.

We attribute this behavior to *equilibrium internalization*: during training, the backbone learns to produce a proposal $\tilde{y}_0$ that already lies close to the fixed point that the solver would otherwise compute through iteration. Thus, the iterative block still shapes the learned representation during training, but at inference time the model has largely internalized the result of that refinement into its initialization. Strikingly, the backbone-only readout at $T = 0$ is already stronger than larger standard Transformers trained without attractor assistance, and matches or exceeds finite-loop baselines that require many test-time recurrences. Concretely, our 140M model at $T = 0$ matches Parcae 140M at $T = 8$; the same pattern holds at 370M and 770M.

Takeaway: Equilibrium Internalization

Attractor Models learn to initialize a prediction near the fixed point, so inference is effectively converged before any explicit refinement step. As a result, they preserve the benefits of iterative computation during training while avoiding the need for many sequential iterations at test time, unlike standard looped models whose quality depends on repeated inference-time updates.

4.4. Ablations

For the ablations, we use a 60.3M-parameter version of our model trained on 1B tokens of FineWeb-Edu under the same `nanochat` configuration. All ablation runs share the same data stream, tokenizer, and optimizer; only the ablated

Table 3. Comparison to DEQ at matched parameter count.

Method	Size	Equilibrium	Val. PPL ↓	Avg. Iters ↓	Core ↑
DEQ (Bai et al., 2019)	60.3M	hidden ($z_0 = 0$, sep. head)	42.18	14.6	5.21 ± 0.14
DEQ + tied unemb.	60.3M	hidden ($z_0 = 0$)	38.74	13.9	5.83 ± 0.12
Attractor Model	60.3M	output emb. ($\tilde{y}_0 = \mathcal{T}_{\theta_b}(E(x))$)	**34.05**	**8.4**	$\mathbf{6.74 \pm 0.10}$

component varies.

We isolate the contribution of (i) the location of the equilibrium and its initialization (Table 3), (ii) the additive injection of c (Table 4), (iii) and the one-step backward approximation (Table 5). All ablations train a 60.3M-parameter model on 1B tokens of FineWeb-Edu under the same `nanochat` setup, with all other hyperparameters fixed.

Comparison to DEQ. Table 3 compares Attractor Models against a parameter-matched DEQ (Bai et al., 2019). The DEQ baseline solves for a hidden-state equilibrium initialized from an uninformative state and decodes it with a separate output head. Tying the unembedding closes part of the gap (Val. PPL $42.18 \rightarrow 38.74$), but the largest improvement comes from matching the equilibrium to the design of Attractor Models: the fixed point is placed directly in the tied output-embedding space and the root finder is initialized from the backbone proposal $\tilde{y}_0 = \mathcal{T}_{\theta_b}(E(x))$ ($\rightarrow 34.05$).

This proposal gives the solver a meaningful, input-dependent starting point rather than requiring the recurrent block to construct the representation from scratch. As a result, the Attractor Model reaches the same residual tolerance in $1.7\times$ fewer iterations while also achieving lower perplexity. We view this as evidence of *equilibrium internalization*: the backbone learns to produce an output-side proposal $\tilde{y}_0$ that already lies close to the eventual equilibrium $\tilde{y}^\star$, making the subsequent attractor refinement both easier and faster.

Proposal injection. Table 4 ablates how the backbone proposal $\tilde{y}_0$ is provided to the attractor. We use *proposal injection* to denote how $\tilde{y}_0$ enters the recurrent refinement map, and *persistent injection* to denote providing $\tilde{y}_0$ at every refinement step rather than only through the initial state.

When the proposal is used only for initialization, $\tilde{y}_{t=0} = \tilde{y}_0$ and $\tilde{y}_{t+1} = \mathcal{T}_{\theta_a}(\tilde{y}_t)$, the recurrent update no longer depends on the input after the first state. Consequently, the fixed point can become proposal-independent, and only 12.4% of validation tokens converge within $T_{\max}$. Persistent injection by concatenation, $\tilde{y}_{t+1} = \mathcal{T}_{\theta_a}([\tilde{y}_t; \tilde{y}_0])$, restores proposal-dependence and recovers much of the quality, but makes the refinement problem harder: convergence is slower (11.2 vs. 8.4 average iterations) and perplexity is worse (36.81 vs. 34.05). Additive proposal injection, $\tilde{y}_{t+1} = \mathcal{T}_{\theta_a}(\tilde{y}_t, \tilde{y}_0)$, provides the backbone proposal at every step while keeping the refinement map simple and well-conditioned, yielding the best results across all three metrics.

Table 4. Effect of input injection on convergence and quality.

Injection of c	Val. PPL ↓	Avg. Iters ↓	Converge ↑
Initial only	51.92	$T_{\max}$	12.4%
Concatenation	36.81	11.2	88.6%
Additive (ours):	**34.05**	**8.4**	**99.7%**

Backward pass. In table 5, we compare the one-step backward approximation against full implicit differentiation (Anderson on the linear system $(I - J_{\tilde{y}}^\top)u = v$) and the phantom gradient (Geng et al., 2022) unroll. Full IFT improves PPL by only 0.14 while increasing peak training memory by $4.8\times$ and step time by $2.7\times$; phantom gradient lies in between. The one-step approximation allows relatively cheap training cost while maintaining nearly all of the original performance of a full backward gradient computation.

Table 5. Ablation for running the full IFT gradients, phantom gradients, and the one-step.

Backward Pass	Val. PPL ↓	Train Mem. ↓	Step Time ↓	Δ vs. Full
Full IFT (Anderson on u)	33.91	$4.8\times$	$2.7\times$	—
Phantom gradient (k=3)	34.02	$1.8\times$	$1.4\times$	+0.11
Ours: one-step ($u \approx v$)	**34.05**	$\mathbf{1.0\times}$	$\mathbf{1.0\times}$	+0.14

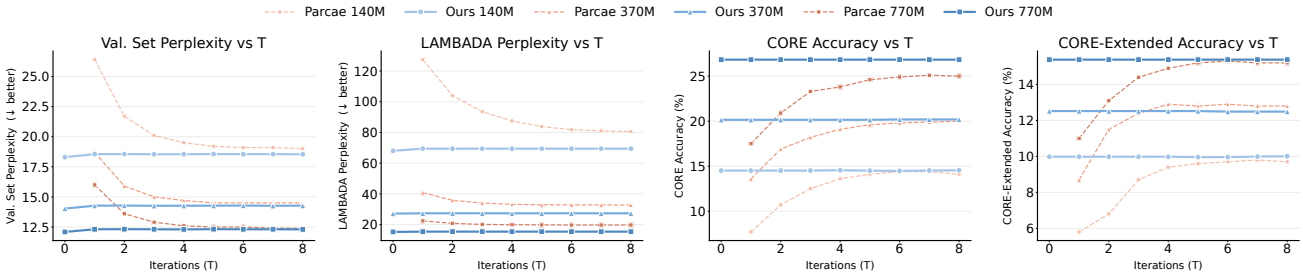

Figure 5. Validation perplexity on FineWeb-Edu, Lambada, CORE accuracy, and CORE-Extended accuracy as a function of test-time iterations T, at three scales. Parcae's quality improves monotonically up to $T \approx \mu_{\text{rec}} = 8$. Our method is essentially converged at $T = 1$ (and at $T = 0$ for the 770M model), because the solver is initialized at $\tilde{y}_0$, which is itself a coherent prediction.

References

Anderson, D. G. Iterative procedures for nonlinear integral equations. *Journal of the ACM*, 12(4):547–560, 1965. doi: 10.1145/321296.321305.

Anthropic. The claude 3 model family: Opus, Sonnet, Haiku. Model card, 2024. URL https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.

Bae, S., Fisch, A., Harutyunyan, H., Ji, Z., Kim, S., and Schuster, T. Relaxed recursive transformers: Effective parameter sharing with layer-wise lora, 2025a. URL https://arxiv.org/abs/2410.20672.

Bae, S., Kim, Y., Bayat, R., Kim, S., Ha, J., Schuster, T., Fisch, A., Harutyunyan, H., Ji, Z., Courville, A., and Yun, S.-Y. Mixture-of-recursions: Learning dynamic recursive depths for adaptive token-level computation, 2025b. URL https://arxiv.org/abs/2507.10524.

Bai, S., Kolter, J. Z., and Koltun, V. Deep equilibrium models. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, 2019.

Bai, S., Koltun, V., and Kolter, J. Z. Multiscale deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/3812f9a59b634c2a9c574610eaba5bed-Abstract.html.

Bai, S., Koltun, V., and Kolter, J. Z. Stabilizing equilibrium models by jacobian regularization, 2021. URL https://arxiv.org/abs/2106.14342.

Blayney, H., Álvaro Arroyo, Obando-Ceron, J., Castro, P. S., Courville, A., Bronstein, M. M., and Dong, X. A mechanistic analysis of looped reasoning language models, 2026. URL https://arxiv.org/abs/2604.11791.

Cameron, C., Wang, W., Ivanov, N., Bhattacharyya, A., Chételat, D., and Zhang, Y. One step forward and k steps back: Better reasoning with denoising recursion models, 2026. URL https://arxiv.org/abs/2604.18839.

Chen, Y., Shang, J., Zhang, Z., Xie, Y., Sheng, J., Liu, T., Wang, S., Sun, Y., Wu, H., and Wang, H. Inner thinking transformer: Leveraging dynamic depth scaling to foster adaptive internal thinking, 2025. URL https://arxiv.org/abs/2502.13842.

Dabre, R. and Fujita, A. Recurrent stacking of layers for compact neural machine translation models, 2018. URL https://arxiv.org/abs/1807.05353.

Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Łukasz Kaiser. Universal transformers, 2019. URL https://arxiv.org/abs/1807.03819.

Deng, J., Pang, L., Wei, Z., Xu, S., Duan, Z., Xu, K., Song, Y., Shen, H., and Cheng, X. Latent reasoning in LLMs as a vocabulary-space superposition. *arXiv preprint arXiv:2510.15522*, 2025.

Deng, J., Wei, Z., Pang, L., Wu, J., Xu, S., Duan, Z., and Shen, H. Latent-GRPO: Group relative policy optimization for latent reasoning. *arXiv preprint arXiv:2604.27998*, 2026.

Fan, Y., Du, Y., Ramchandran, K., and Lee, K. Looped transformers for length generalization, 2025. URL https://arxiv.org/abs/2409.15647.

Fleuret, F. The free transformer, 2025. URL https://arxiv.org/abs/2510.17558.

Fu, T., You, Y., Chen, Z., Dai, G., Yang, H., and Wang, Y. Think-at-hard: Selective latent iterations to improve reasoning language models, 2026. URL https://arxiv.org/abs/2511.08577.

Fung, S. W., Heaton, H., Li, Q., McKenzie, D., Osher, S., and Yin, W. Jfb: Jacobian-free backpropagation for implicit networks, 2021. URL https://arxiv.org/abs/2103.12803.

Gatmiry, K., Saunshi, N., Reddi, S. J., Jegelka, S., and Kumar, S. Can looped transformers learn to implement multi-step gradient descent for in-context learning? In *International Conference on Machine Learning*, pp. 15130–15152. PMLR, 2024a.

Gatmiry, K., Saunshi, N., Reddi, S. J., Jegelka, S., and Kumar, S. On the role of depth and looping for in-context learning with task diversity, 2024b. URL https://arxiv.org/abs/2410.21698.

Geiping, J., McLeish, S., Jain, N., Kirchenbauer, J., Singh, S., Bartoldson, B. R., Kailkhura, B., Bhatele, A., and Goldstein, T. Scaling up test-time compute with latent reasoning: A recurrent depth approach, 2025a. URL https://arxiv.org/abs/2502.05171.

Geiping, J., Yang, X., and Su, G. Efficient parallel samplers for recurrent-depth models and their connection to diffusion language models, 2025b. URL https://arxiv.org/abs/2510.14961.

Gemini Team. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

Geng, Z., Zhang, X.-Y., Bai, S., Wang, Y., and Lin, Z. On training implicit models, 2022. URL https://arxiv.org/abs/2111.05177.

Ghugare, R., Bortkiewicz, M., Ziarko, A., and Eysenbach, B. On the role of iterative computation in reinforcement learning, 2026. URL https://arxiv.org/abs/2602.05999.

Giannou, A., Rajput, S., Sohn, J.-y., Lee, K., Lee, J. D., and Papailiopoulos, D. Looped transformers as programmable computers. In *International Conference on Machine Learning*, pp. 11398–11442. PMLR, 2023a.

Giannou, A., Rajput, S., yong Sohn, J., Lee, K., Lee, J. D., and Papailiopoulos, D. Looped transformers as programmable computers, 2023b. URL https://arxiv.org/abs/2301.13196.

Goyal, S., Agrawal, S., Anil, G. G., Jain, P., Paul, S., and Kusupati, A. Elt: Elastic looped transformers for visual generation, 2026. URL https://arxiv.org/abs/2604.09168.

Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., et al. The Llama 3 herd of models. In *Neural Information Processing Systems*, 2024.

Graves, A. Adaptive computation time for recurrent neural networks, 2017. URL https://arxiv.org/abs/1603.08983.

Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z. F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Xue, B., Wang, B., Wu, B., Feng, B., Lu, C., Zhao, C., Deng, C., Ruan, C., Dai, D., Chen, D., Ji, D., Li, E., Lin, F., Dai, F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Xu, H., Ding, H., Gao, H., Qu, H., Li, H., Guo, J., Li, J., Chen, J., Yuan, J., Tu, J., Qiu, J., Li, J., Cai, J. L., Ni, J., Liang, J., Chen, J., Dong, K., Hu, K., You, K., Gao, K., Guan, K., Huang, K., Yu, K., Wang, L., Zhang, L., Zhao, L., Wang, L., Zhang, L., Xu, L., Xia, L., Zhang, M., Zhang, M., Tang, M., Zhou, M., Li, M., Wang, M., Li, M., Tian, N., Huang, P., Zhang, P., Wang, Q., Chen, Q., Du, Q., Ge, R., Zhang, R., Pan, R., Wang, R., Chen, R. J., Jin, R. L., Chen, R., Lu, S., Zhou, S., Chen, S., Ye, S., Wang, S., Yu, S., Zhou, S., Pan, S., Li, S. S., Zhou, S., Wu, S., Yun, T., Pei, T., Sun, T., Wang, T., Zeng, W., Liu, W., Liang, W., Gao, W., Yu, W., Zhang, W., Xiao, W. L., An, W., Liu, X., Wang, X., Chen, X., Nie, X., Cheng, X., Liu, X., Xie, X., Liu, X., Yang, X., Li, X., Su, X., Lin, X., Li, X. Q., Jin, X., Shen, X., Chen, X., Sun, X., Wang, X., Song, X., Zhou, X., Wang, X., Shan, X., Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhang, Y., Xu, Y., Li, Y., Zhao, Y., Sun, Y., Wang, Y., Yu, Y., Zhang, Y., Shi, Y., Xiong, Y., He, Y., Piao, Y., Wang, Y., Tan, Y., Ma, Y., Liu, Y., Guo, Y., Ou, Y., Wang, Y., Gong, Y., Zou, Y., He, Y., Xiong, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Zhu, Y. X., Huang, Y., Li, Y., Zheng, Y., Zhu, Y., Ma, Y., Tang, Y., Zha, Y., Yan, Y., Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Xie, Z., Zhang, Z., Hao, Z., Ma, Z., Yan, Z., Wu, Z., Gu, Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Pan, Z., Huang, Z., Xu, Z., Zhang, Z., and Zhang, Z. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645 (8081):633–638, 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-09422-z. URL http://dx.doi.org/10.1038/s41586-025-09422-z.

Han, S.-I. Hierarchical vs. flat iteration in shared-weight transformers, 2026. URL https://arxiv.org/abs/2604.14442.

Hao, S., Sukhbaatar, S., Su, D., Li, X., Hu, Z., Weston, J., and Tian, Y. Training large language models to reason in a continuous latent space, 2025. URL https://arxiv.org/abs/2412.06769.

Hu, E. S., Ahn, K., Liu, Q., Xu, H., Tomar, M., Langford, A., Teoh, J., Xu, B., Yan, D., Jayaraman, D., Lamb, A., and Langford, J. The belief state transformer, 2025. URL https://arxiv.org/abs/2410.23506.

Huang, J., Wang, Z., and Lee, J. D. Transformers learn to implement multi-step gradient descent with chain of thought, 2025. URL https://arxiv.org/abs/2502.21212.

Inan, H., Khosravi, K., and Socher, R. Tying word vectors and word classifiers: A loss framework for language modeling, 2017. URL https://arxiv.org/abs/1611.01462.

Jeddi, A., Ciccone, M., and Taati, B. Loopformer: Elastic-depth looped transformers for latent reasoning via shortcut modulation, 2026. URL https://arxiv.org/abs/2602.11451.

Jolicoeur-Martineau, A. Less is more: Recursive reasoning with tiny networks, 2025. URL https://arxiv.org/abs/2510.04871.

Karpathy, A. nanochat: The best chatgpt that $100 can buy, 2025. URL https://github.com/karpathy/nanochat.

Knupp, J., Metzen, J. H., Bohn, J., Groh, G., and Kersting, K. Depth-Recurrent Attention Mixtures: Giving latent reasoning the attention it deserves. *arXiv preprint arXiv:2601.21582*, 2026a.

Knupp, J., Metzen, J. H., Bohn, J., Groh, G., and Kersting, K. Depth-recurrent attention mixtures: Giving latent reasoning the attention it deserves, 2026b. URL https://arxiv.org/abs/2601.21582.

Kohli, H., Parthasarathy, S., Sun, H., and Yao, Y. Loop, think, & generalize: Implicit reasoning in recurrent-depth transformers, 2026. URL https://arxiv.org/abs/2604.07822.

Komisarczyk, M., Mathur, S., Kraus, M., Natarajan, S., and Kersting, K. Recursive inference machines for neural reasoning, 2026. URL https://arxiv.org/abs/2603.05234.

Krantz, S. G. and Parks, H. R. *The Implicit Function Theorem: History, Theory, and Applications*. Birkhäuser, Boston, MA, 2002a. ISBN 978-0-8176-4285-3. doi: 10.1007/978-1-4612-0059-8.

Krantz, S. G. and Parks, H. R. *The implicit function theorem: history, theory, and applications*, volume 202. Springer, 2002b.

Labovich, A. Stability and generalization in looped transformers, 2026a. URL https://arxiv.org/abs/2604.15259.

Labovich, A. Stability and generalization in looped transformers, 2026b. URL https://arxiv.org/abs/2604.15259.

Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. Albert: A lite bert for self-supervised learning of language representations, 2020. URL https://arxiv.org/abs/1909.11942.

McLeish, S. M., Li, A., Kirchenbauer, J., Kalra, D. S., Bartoldson, B. R., Kailkhura, B., Schwarzschild, A., Geiping, J., Goldblum, M., and Goldstein, T. Teaching pretrained language models to think deeper with retrofitted recurrence. In *NeurIPS 2025 Workshop on Efficient Reasoning*.

Merrill, W. and Sabharwal, A. A little depth goes a long way: The expressive power of log-depth transformers, 2025. URL https://arxiv.org/abs/2503.03961.

Merrill, W. and Sabharwal, A. A little depth goes a long way: The expressive power of log-depth transformers. *Advances in Neural Information Processing Systems*, 38: 95315–95339, 2026.

Mohtashami, A., Pagliardini, M., and Jaggi, M. CoTFormer: More tokens with attention make up for less depth. In *Workshop on Advancing Neural Network Training: Computational Efficiency, Scalability, and Resource Optimization (WANT@ NeurIPS 2023)*, 2023.

Moosa, I. M., Lohit, S., Wang, Y., Chatterjee, M., and Yin, W. Understanding dynamic compute allocation in recurrent transformers, 2026. URL https://arxiv.org/abs/2602.08864.

OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Özeren, E. and Aßenmacher, M. Reinforcement learning for latent-space thinking in LLMs. *arXiv preprint arXiv:2512.11816*, 2025.

Paperno, D., Kruszewski, G., Lazaridou, A., Pham, N. Q., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1525–1534, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/P16-1144. URL https://aclanthology.org/P16-1144/.

Pappone, F., Crisostomi, D., and Rodolà, E. Two-scale latent dynamics for recurrent-depth transformers, 2025. URL https://arxiv.org/abs/2509.23314.

Penedo, G., Kydlíček, H., allal, L. B., Lozhkov, A., Mitchell, M., Raffel, C., Werra, L. V., and Wolf, T. The fineweb datasets: Decanting the web for the finest text data at scale, 2024. URL https://arxiv.org/abs/2406.17557.

Prairie, H., Novack, Z., Berg-Kirkpatrick, T., and Fu, D. Y. Parcae: Scaling laws for stable looped language models, 2026. URL https://arxiv.org/abs/2604.12946.

Press, O. and Wolf, L. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 157–163, 2017.

Rauba, P., Fanconi, C., and van der Schaar, M. Tiny autoregressive recursive models, 2026. URL https://arxiv.org/abs/2603.08082.

Rizvi-Martel, M., Rabusseau, G., and Mosbach, M. The illusion of superposition? A principled analysis of latent thinking in language models. *arXiv preprint arXiv:2604.06374*, 2026.

Sapunov, G. Universal transformers need memory: Depth-state trade-offs in adaptive recursive reasoning, 2026a. URL https://arxiv.org/abs/2604.21999.

Sapunov, G. Universal transformers need memory: Depth-state trade-offs in adaptive recursive reasoning, 2026b. URL https://arxiv.org/abs/2604.21999.

Saunshi, N., Dikkala, N., Li, Z., Kumar, S., and Reddi, S. J. Reasoning with latent thoughts: On the power of looped transformers, 2025a. URL https://arxiv.org/abs/2502.17416.

Saunshi, N., Dikkala, N., Li, Z., Kumar, S., and Reddi, S. J. Reasoning with latent thoughts: On the power of looped transformers, 2025b. URL https://arxiv.org/abs/2502.17416.

Song, S., Li, H., Wang, Z., Zeng, B., Song, F., Wang, Y., Xu, Z. J., He, Z., and Lin, Z. AdaPonderLM: Gated pondering language models with token-wise adaptive depth. *arXiv preprint arXiv:2603.01914*, 2026a.

Song, S., Li, H., Wang, Z., Zeng, B., Song, F., Wang, Y., Xu, Z. J., He, Z., and Lin, Z. Adaponderlm: Gated pondering language models with token-wise adaptive depth, 2026b. URL https://arxiv.org/abs/2603.01914.

Takase, S. and Kiyono, S. Lessons on parameter sharing across layers in transformers, 2023. URL https://arxiv.org/abs/2104.06022.

Tang, G., Jiang, S., Chang, H., Chen, N., Li, Y., Fan, H., Li, J., Liu, M., and Qin, B. Looprpt: Reinforcement pre-training for looped language models, 2026. URL https://arxiv.org/abs/2603.19714.

Teoh, J., Tomar, M., Ahn, K., Hu, E. S., Sharma, P., Islam, R., Lamb, A., and Langford, J. Next-latent prediction transformers learn compact world models, 2025. URL https://arxiv.org/abs/2511.05963.

Tur, Y., Naghiyev, J., Fang, H., Tsai, W.-C., Duan, J., Fox, D., and Krishna, R. Recurrent-depth vla: Implicit test-time compute scaling of vision-language-action models via latent iterative reasoning, 2026. URL https://arxiv.org/abs/2602.07845.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems 30*, pp. 5998–6008, 2017.

Wang, G., Li, J., Sun, Y., Chen, X., Liu, C., Wu, Y., Lu, M., Song, S., and Yadkori, Y. A. Hierarchical reasoning model, 2025. URL https://arxiv.org/abs/2506.21734.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., and Zhou, D. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL https://arxiv.org/abs/2201.11903.

Wei, X., Liu, X., Zang, Y., Dong, X., Cao, Y., Wang, J., Qiu, X., and Lin, D. Sim-cot: Supervised implicit chain-of-thought, 2025. URL https://arxiv.org/abs/2509.20317.

Williams, J. and Tureci, E. Prioritize the process, not just the outcome: Rewarding latent thought trajectories improves reasoning in looped language models, 2026. URL https://arxiv.org/abs/2602.10520.

Wu, B., Chen, M., Luo, X., Yan, S., Yu, Q., Xia, F., Zhang, T., Zhan, H., Zhong, Z., Zhou, X., Qiao, S., and Bin, X. Parallel loop transformer for efficient test-time computation scaling, 2025. URL https://arxiv.org/abs/2510.24824.

Xu, K. and Sato, I. On expressive power of looped transformers: Theoretical analysis and enhancement via timestep encoding, 2025a. URL https://arxiv.org/abs/2410.01405.

Xu, K. and Sato, I. On expressive power of looped transformers: Theoretical analysis and enhancement via timestep encoding, 2025b. URL https://arxiv.org/abs/2410.01405.

Xu, K. and Sato, I. A formal comparison between chain of thought and latent thought, 2026. URL https://arxiv.org/abs/2509.25239.

Yang, L., Lee, K., Nowak, R., and Papailiopoulos, D. Looped transformers are better at learning learning algorithms, 2024a. URL https://arxiv.org/abs/2311.12424.

Yang, L., Lee, K., Nowak, R., and Papailiopoulos, D. Looped transformers are better at learning learning algorithms, 2024b. URL https://arxiv.org/abs/2311.12424.

Yang, L., Lee, K., Nowak, R., and Papailiopoulos, D. Looped transformers are better at learning learning algorithms, 2024c. URL https://arxiv.org/abs/2311.12424.

Zeitoun, A., Torroba-Hennigen, L., and Kim, Y. Hyperloop transformers. *arXiv preprint arXiv:2604.21254*, 2026a.

Zeitoun, A., Torroba-Hennigen, L., and Kim, Y. Hyperloop transformers, 2026b. URL https://arxiv.org/abs/2604.21254.

Zeng, B., Song, S., Huang, S., Wang, Y., Li, H., He, Z., Wang, X., Li, Z., and Lin, Z. Ponderlm: Pretraining language models to ponder in continuous space, 2026. URL https://arxiv.org/abs/2505.20674.

Zhang, X., Wu, H., He, G., Shen, J., Lyu, B., and Zhu, Z. MoDr: Mixture-of-depth-recurrent transformers for test-time reasoning. In *The Fourteenth International Conference on Learning Representations*.

Zhang, X., Wu, H., He, G., Shen, J., Lyu, B., and Zhu, Z. Modr: Mixture-of-depth-recurrent transformers for test-time reasoning. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=9Pba4rcQbE.

Zheng, Y., Wang, X., Ren, L., and Wei, C. ChainGPT: Dual-reasoning model with recurrent depth and multi-rank state updates. In *The Fourteenth International Conference on Learning Representations*.

Zheng, Y., Wang, X., Ren, L., and Wei, C. ChainGPT: Dual-reasoning model with recurrent depth and multi-rank state updates. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=kdZbxizwGK.

Zhou, C., Yang, C., Hu, Y., Wang, C., Zhang, C., Zhang, M., Mackey, L., Jaakkola, T., Bates, S., and Zhang, D. Coevolutionary continuous discrete diffusion: Make your diffusion language model a latent reasoner. *arXiv preprint arXiv:2510.03206*, 2025.

Zhu, H., Hao, S., Hu, Z., Jiao, J., Russell, S., and Tian, Y. Emergence of superposition: Unveiling the training dynamics of chain of continuous thought. *arXiv preprint arXiv:2509.23365*, 2025a.

Zhu, H., Hao, S., Hu, Z., Jiao, J., Russell, S., and Tian, Y. Reasoning by superposition: A theoretical perspective on chain of continuous thought, 2025b. URL https://arxiv.org/abs/2505.12514.

Zhu, H., Hao, S., Hu, Z., Jiao, J., Russell, S. J., and Tian, Y. Reasoning by superposition: A theoretical perspective on chain of continuous thought. *Advances in Neural Information Processing Systems*, 38:79931–79963, 2026.

Zhu, R.-J., Wang, Z., Hua, K., Zhang, T., Li, Z., Que, H., Wei, B., Wen, Z., Yin, F., Xing, H., Li, L., Shi, J., Ma, K., Li, S., Kergan, T., Smith, A., Qu, X., Hui, M., Wu, B., Min, Q., Huang, H., Zhou, X., Ye, W., Liu, J., Yang, J., Shi, Y., Lin, C., Zhao, E., Cai, T., Zhang, G., Huang, W., Bengio, Y., and Eshraghian, J. Scaling latent reasoning via looped language models, 2025c. URL https://arxiv.org/abs/2510.25741.

Łukasz Kaiser and Sutskever, I. Neural gpus learn algorithms, 2016. URL https://arxiv.org/abs/1511.08228.

A. Related Work

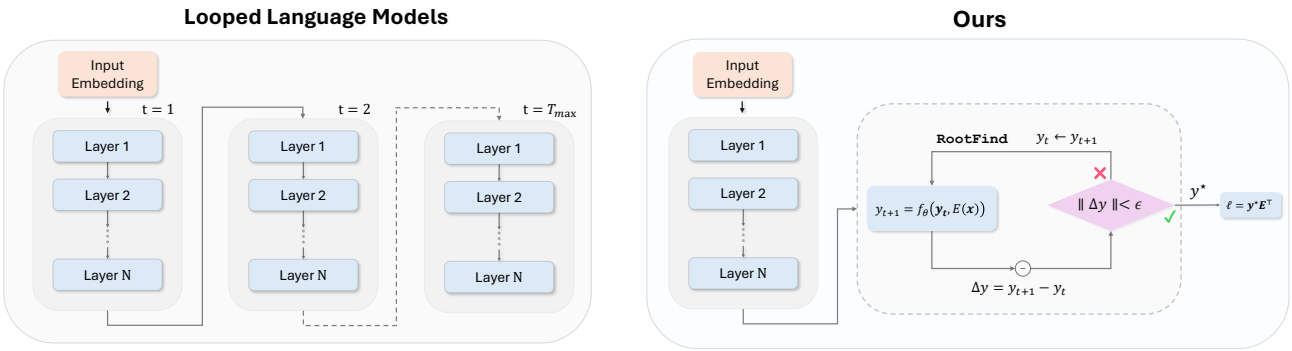

Figure 7. Comparison of standard looped language models vs. our architecture

Looped and recurrent language models. Weight-tied recurrence has re-emerged as an alternative to deeper feed-forward stacks (Press & Wolf, 2017; Inan et al., 2017; Łukasz Kaiser & Sutskever, 2016; Lan et al., 2020; Dabre & Fujita, 2018; Takase & Kiyono, 2023; Bae et al., 2025a). Universal Transformers share a single block across depth (Dehghani et al., 2019; Sapunov, 2026a;b), while looped and recurrent language models iterate a recurrent update on a hidden state to enable latent reasoning (Zhu et al., 2025c; Prairie et al., 2026; Geiping et al., 2025a; Bae et al., 2025b; Zheng et al.; Zhang et al.; Song et al., 2026a; Knupp et al., 2026a; Zeitoun et al., 2026a; McLeish et al.; Zheng et al., 2026; Zhang et al., 2026; Song et al., 2026b; Moosa et al., 2026; Knupp et al., 2026b; Wu et al., 2025; Zeitoun et al., 2026b; Pappone et al., 2025; Rauba et al., 2026; Zeng et al., 2026; Tur et al., 2026; Komisarczyk et al., 2026; Ghugare et al., 2026; Cameron et al., 2026; Williams & Tureci, 2026; Han, 2026; Goyal et al., 2026; Tang et al., 2026). Latent (implicit) reasoning methods such as Coconut (Hao et al., 2025) and related methods (Teoh et al., 2025; Chen et al., 2025; Xu & Sato, 2026; Jeddi et al., 2026; Fu et al., 2026; Zhou et al., 2025; Geiping et al., 2025b; Hu et al., 2025; Fleuret, 2025) perform reasoning in continuous representation space rather than through chain-of-thought tokens. These approaches unroll the loop for a fixed number of steps at training time, causing training memory to grow linearly in depth, coupling inference iterations to training, and often degrading quality when iterations are extended at test time (Zhu et al., 2025c). From theoretical and mechanistic analysis perspectives, looped and continuous thinking models have shown to offer benefits over non-recurrent models (Yang et al., 2024a; Labovich, 2026b; Saunshi et al., 2025a; Xu & Sato, 2025a; Saunshi et al., 2025b; Giannou et al., 2023b; Yang et al., 2024c; Gatmiry et al., 2024b; Huang et al., 2025; Xu & Sato, 2025b; Merrill & Sabharwal, 2025; Labovich, 2026a; Zhu et al., 2025b;a; Gatmiry et al., 2024a; Merrill & Sabharwal, 2026). Mechanistic analysis shows that the recurrent updates of looped LMs in fact converge to a fixed point for the vast majority of tokens, directly motivating our formulation (Blayney et al., 2026).

Implicit fixed-point models. Deep Equilibrium Models (DEQs) replace finite unrolling with a fixed-point equation $z^\star = f_\theta(z^\star, x)$ and train through it via implicit differentiation, decoupling effective depth from training memory (Bai et al., 2019). Solvers typically use Anderson acceleration (Anderson, 1965), with backward passes ranging from full implicit

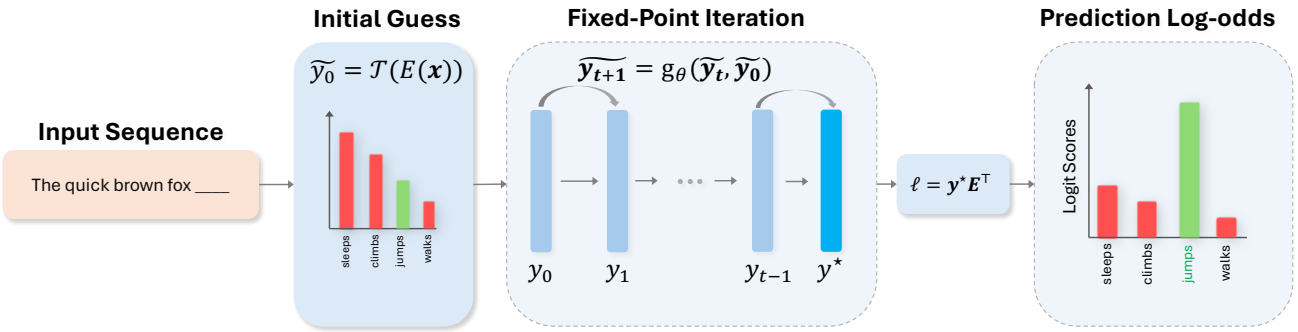

Figure 8. Overview of our method. The backbone maps the input embeddings to an output-side proposal $\tilde{y}_0 = \mathcal{T}_{\theta_b}(E(x))$. The attractor module then refines this proposal through the proposal conditioned iteration $\tilde{y}_{t+1} = \mathcal{T}_{\theta_a}(\tilde{y}_t, \tilde{y}_0)$ until convergence to an equilibrium $\tilde{y}^\star$ or a maximum number of solver steps. The equilibrium is decoded with the tied unembedding $\tilde{y}^\star E^\top$.

differentiation to cheap surrogates such as one-step (Fung et al., 2021) and phantom (Geng et al., 2022) gradients. Standard DEQs solve for a hidden state initialized from zero and decode through a separate output head. In Attractor Models, the equilibrium instead lives in the tied output-embedding space and is warm-started from a meaningful backbone prediction; we find this structure essential for stable training at language-modeling scale and show it gives rise to equilibrium internalization (Section 4.3).

Tiny recursive reasoners. On small-data algorithmic benchmarks such as Sudoku and maze solving, hierarchical and tiny recursive networks (HRM, TRM) (Wang et al., 2025; Jolicoeur-Martineau, 2025) achieve strong accuracy with few parameters, but exhibit a "less is more" behavior where performance collapses as model size grows (Jolicoeur-Martineau, 2025). Attractor Models retain the iterative-refinement benefit of these architectures while scaling cleanly with parameter count).

Positioning. Looped LMs couple three quantities through unrolled BPTT: inference depth, training depth, and training memory (Zhu et al., 2025c). Attractor Models decouple all three: the equilibrium is defined by a residual equation in output-embedding space, the number of solver evaluations is chosen adaptively by tolerance ε, and training memory in the recurrent block is constant in effective depth. Empirically, this combines the strengths of feed-forward and recurrent models—better perplexity than parameter-matched Transformers, 25–31% lower training FLOPs and constant memory relative to looped LMs, and clean scaling on hard reasoning where specialized recursive networks collapse.

B. Theory of Looped and Attractor Models

B.1. Well-posedness and the implicit gradient

Fix an input x and parameters θ. Throughout this section, we write $F(y) \triangleq f_\theta(y, E(x))$ for brevity, so that the fixed-point equation $y^\star = f_\theta(y^\star, E(x))$ becomes $y^\star = F(y^\star)$. Let $J_F(y) \triangleq \partial F/\partial y$ denote the Jacobian of F with respect to its state input. We work in a generic norm $\|\cdot\|$ on $\mathbb{R}^{L\times d}$ and its induced operator norm; for concreteness, can take both to be Frobenius and spectral, respectively. Assumptions are borrowed from (Bai et al., 2020; Krantz & Parks, 2002a).

Assumption B.1 (Local contraction)**.** There exist a point $\bar{y} \in \mathbb{R}^{L\times d}$, a radius $r > 0$, and a constant $L \in [0, 1)$ such that

1. F maps the closed ball $B_r(\bar{y}) \triangleq \{y : \|y - \bar{y}\| \le r\}$ into itself, and
2. F is L-Lipschitz on $B_r(\bar{y})$:

$$\|F(y) - F(y')\| \;\le\; L\,\|y - y'\| \qquad \text{for all } y, y' \in B_r(\bar{y}).$$

When F is continuously differentiable, a sufficient form of (ii) is $\sup_{y\in B_r(\bar{y})} \|J_F(y)\| \le L$.

Theorem B.2 (Well-posedness and implicit gradient)**.** *Under Assumption B.1, the following hold.*

1. Existence and uniqueness: *There is a unique $y^\star \in B_r(\bar{y})$ with $y^\star = F(y^\star)$.*
2. The picard iteration converges linearly: *For any $y_0 \in B_r(\bar{y})$, the iterates $y_{k+1} = F(y_k)$ remain in $B_r(\bar{y})$ and satisfy*

$$\|y_k - y^\star\| \;\le\; L^k\,\|y_0 - y^\star\|. \tag{7}$$

3. Validity of the implicit gradient: *If F is continuously differentiable on $B_r(\bar{y})$, then $I - J_F(y^\star)$ is invertible, $y^\star$ depends continuously differentiably on θ in a neighborhood of the current parameters, and*

$$\frac{\partial y^\star}{\partial \theta} \;=\; \big(I - J_F(y^\star)\big)^{-1}\,\frac{\partial f_\theta}{\partial \theta}\bigg|_{y^\star}. \tag{8}$$

Proof. Parts (i) and (ii) follow from the Banach fixed-point theorem applied to the contraction F on the closed (hence complete) ball $B_r(\bar{y}) \subset (\mathbb{R}^{L\times d}, \|\cdot\|)$.

For (iii), Lipschitz continuity with constant L gives $\|J_F(y^\star)\| \le L < 1$, so the spectral radius satisfies $\rho(J_F(y^\star)) \le \|J_F(y^\star)\| < 1$. The Neumann series $\sum_{k\ge 0} J_F(y^\star)^k$ therefore converges to $(I - J_F(y^\star))^{-1}$, and in particular $I - J_F(y^\star)$ is invertible. Define $G(y, \theta) \triangleq f_\theta(y, E(x)) - y$, which is C^1 in both arguments and satisfies $G(y^\star, \theta) = 0$ with

$$\left.\frac{\partial G}{\partial y}\right|_{y^\star} \;=\; J_F(y^\star) - I, \tag{9}$$

which is invertible. The implicit function theorem applied to $G = 0$ yields a unique C^1 map $\theta \mapsto y^\star(\theta)$ in a neighborhood of the current parameters, with

$$\frac{\partial y^\star}{\partial \theta} \;=\; -\left(\frac{\partial G}{\partial y}\right)^{-1} \frac{\partial G}{\partial \theta} \;=\; \big(I - J_F(y^\star)\big)^{-1} \left.\frac{\partial f_\theta}{\partial \theta}\right|_{y^\star}. \tag{10}$$

□

Equation (8) is exactly the gradient computed in the backward pass of Section 3. The linear solve is convergent for the same reason: since $\|J_F^\top(y^\star)\| = \|J_F(y^\star)\| \le L < 1$, the Neumann iteration $u \leftarrow J_F^\top(y^\star)\, u + v$ converges geometrically with rate L to the unique solution $u = (I - J_F^\top(y^\star))^{-1} v$.

B.2. Looped language models are fixed-point iterators

A LoopLM (Zhu et al., 2025c) of depth T with the same weight-tied block f_θ and warm start $y_0 = E(x)$ produces

$$y_T^{\text{loop}} \;\triangleq\; F^{(T)}\big(E(x)\big) \;=\; \underbrace{F \circ F \circ \cdots \circ F}_{T \text{ times}}\big(E(x)\big). \tag{11}$$

This is exactly T Picard iterations of our residual $g_\theta(\cdot, x)$ warm-started from the input embedding. See the following corollary:

Corollary B.3 (LoopLM as a truncated approximation)**.** *Suppose Assumption B.1 holds and the warm start $E(x) \in B_r(\bar{y})$. Then for every $T \ge 0$,*

$$\|y_T^{\text{loop}} - y^\star\| \;\le\; L^T \, \|E(x) - y^\star\|. \tag{12}$$

In particular, $y_T^{\text{loop}} \to y^\star$ as $T \to \infty$, and the discrepancy between a LoopLM of depth T and our model decays geometrically in T.

Proof. Apply Theorem B.2(ii) with $y_0 = E(x)$. □

Two consequences are worth highlighting:

The limit: Our fixed-point model is the $T \to \infty$ limit of LoopLM with shared parameters. Training a LoopLM at any finite depth T implicitly trains an approximation to the same equilibrium $y^\star$, with the approximation error controlled by (12).

The iteration count at inference: Picard iteration converges at the linear rate L. Our forward pass uses Anderson acceleration, which under standard regularity conditions converges *superlinearly* near $y^\star$ (Anderson, 1965). To reach a target residual tolerance ε, the root finder therefore typically requires fewer evaluations of f_θ than the unroll depth T a LoopLM would need for a comparable residual. The exact crossover depends on L and on the solver hyperparameters.

Caveat: Assumption B.1 is local and is not guaranteed to hold for an arbitrarily trained transformer block. In our experiments, we generally notice that $\|J_F(y^\star)\| \le L < 1$, the authors of DEQ also find that using a regularizer on the Jacobian to enforce this helps (Bai et al., 2021).

B.3. Where looped language models fall short

B.3.1. Attractor Models

The initialization $\tilde{y}_0$ and the fixed point $\tilde{y}^\star$ live in the same tied output-embedding space, so either can be decoded by the unembedding $E^\top$. The loss depends only on $\tilde{y}^\star$, and the implicit gradient with respect to the backbone parameters is

$$\nabla_{\theta_b} \mathcal{L} \;=\; u^\top J_{\tilde{y}_0} \, \nabla_{\theta_b} \tilde{y}_0, \qquad u = (I - J_{\tilde{y}}^\top)^{-1} v.$$

Thus, the loss may be viewed as a function of the backbone's output embedding: perturbations in $\tilde{y}_0$ affect $\mathcal{L}$ by perturbing $\tilde{y}^\star$ in the same space. The backbone is therefore optimized as a standalone next-token predictor whose embedding is decoded by $E^\top$.

When the backbone has strictly greater capacity than the weight-tied attractor, the joint optimum places the prediction work in the backbone. Consequently, $\tilde{y}_0(x; \theta_b^\star)$ approaches the loss-minimizing embedding. The fixed-point constraint

$$\mathcal{T}_{\theta_a^\star}(\tilde{y}^\star, \tilde{y}_0) = \tilde{y}^\star$$

is then consistent only when $\tilde{y}^\star = \tilde{y}_0$, yielding

$$\mathcal{T}_{\theta_a^\star}(\tilde{y}_0, \tilde{y}_0) = \tilde{y}_0.$$

B.3.2. Standard looped LMs

In standard looped language models, the latent initialization h_0 is uninformative, typically zero or noise, and occupies the role of a hidden state rather than an output embedding. Decoding $h_0 E^\top$ is therefore not meaningful.

Gradients instead flow by backpropagation through time along the unrolled trajectory $(h_0, \ldots, h_T)$. There is no term in this gradient that drives h_0 toward the loss-bearing final state h_T in embedding distance: the initialization is fixed by the architecture and is not produced by a separately trained predictor. □

B.3.3. Implicit-Gradient Barrier

The implicit gradient

$$\nabla_\theta \mathcal{L}_\infty = v^\top (I - J_g^\top)^{-1} \partial_\theta g$$

contains the factor $(I - J_g)^{-1}$. This inverse is undefined when $1 \in \sigma(J_g)$, and its operator norm scales as

$$\|(I - J_g)^{-1}\| \sim \mathrm{dist}(1, \sigma(J_g))^{-1}$$

near such a singularity.

Along any continuous gradient-descent path, the eigenvalues of $J_g(y^\star; \theta)$ vary continuously. Therefore, for the dominant eigenvalue to leave the unit disk through $+1$—the canonical loss-of-contraction route for residual iteration maps—it must approach 1. This forces $\|(I - J_g)^{-1}\| \to \infty$.

Unless v is orthogonal to the offending eigendirection, a codimension-one and hence non-generic condition, the gradient norm diverges. The descent step therefore blows up before the boundary can be crossed, confining θ to the contractive region $\{\rho(J_g) < 1\}$.

For exits through complex eigenvalues, the same conclusion holds under the one-step JFB approximation, since the truncated Neumann series

$$\sum_k (J_g^\top)^k$$

diverges as $\rho(J_g) \to 1$. □

B.3.4. Interpretation

Fixed-loop training has no inherent mechanism that favors contractive iterations. An optimum in which g_θ perturbs the embedding for K steps and only happens to land accurately at step K is a valid fixed-loop solution. These are precisely the solutions that fail when extra loops are run at inference.

Equilibrium training cannot reach such solutions, because the implicit gradient creates a barrier of diverging gradients around non-contractive regions. As a result, the trajectory of θ is confined to the regime in which the solver converges, the one-step gradient is a descent direction, and additional iterations remain stable.

The mechanistic observation (Blayney et al., 2026) that fixed points emerge only late in standard looped training is consistent with this picture: the loss landscape contains a basin near contractive solutions, but only equilibrium training applies pressure toward that basin from the beginning of training.

C. Hyperparamter and Experimental Settings

We use NVIDIA H200 GPUs for all of our experiments.

Table 6. Architecture hyperparameters for each model family and scale.

	Transformer			**Parcae**			**Attractor Model**		
Hyperparameter	140M	370M	770M	140M	370M	770M	140M	370M	770M
d_{model}	768	1024	1280	768	1024	1280	1024	1280	1280
d_{ff}	3072	4096	5120	3072	4096	5120	4096	5120	5120
Attention heads	6	8	10	6	8	10	8	10	10
Total layers	6	12	18	6	12	18	8	17	37
Prelude	—	—	—	2	4	6	7	15	35
Core (recurrent)	—	—	—	2	4	6	1	2	2
Coda	—	—	—	2	4	6	0	0	0
Mean recurrence T	1	1	1	8	8	8	(solver-determined)		
Sequence length	2048								
Vocab size	32768								
Tie embeddings	✓								
Norm	RMSNorm ($\epsilon = 10^{-5}$)								
Activation	ReLU^2								
QK-norm	✓								
Precision	bf16-mixed								

Table 7. Fixed-point solver hyperparameters for Attractor Model. Parcae and Transformer do not use a solver.

Hyperparameter	**140M**	**370M**	**770M**
Forward solver	Anderson	Anderson	Anderson
Max iterations (fwd)	64	64	32
Min iterations (fwd)	6	8	6
Tolerance (fwd)	3×10^{-4}	2×10^{-4}	1.5×10^{-4}
Anderson m	5	5	5
Anderson β	1.0	1.0	0.96
Backward type	one-step IFT	one-step IFT	Anderson
Max iterations (bwd)	64	64	32
Min iterations (bwd)	6	6	6
Tolerance (bwd)	3×10^{-4}	2×10^{-4}	1.5×10^{-4}
Adjoint grad clip	—	—	2.0
Layer-scale init (γ_0)	0.75	0.75	0.73
$\gamma_{\max}$	0.75	1.0	0.9
FP LR scale	0.5	0.4	0.4
FP weight decay	0.1	0.1	0.1

D. Conclusion and Future Work

In this work, we propose Attractor Models, a new family of architectures that first produce meaningful prediction embeddings and then refine them through an attractor module by solving for a fixed point. This formulation makes recurrent refinement stable, memory-efficient, and adaptive, while avoiding the cost of explicit unrolling and achieving strong results across both large-scale language modeling and hard reasoning with tiny models. In stark contrast to prior looped language models, training Attractor Models gives rise to equilibrium internalization: the model learns to make the very refinement procedure

Table 8. Training hyperparameters. All three model families share the same optimizer, schedule, and data configuration at each scale. Differences are noted where they occur.

Hyperparameter	**140M**	**370M**	**770M**
Training tokens	11.2B	29.6B	61.6B
Dataset	FineWeb-Edu 100B (shuffled)		
Tokenizer	BPE, 32768 vocab		
Optimizer	MuonAdamW		
AdamW LR	8×10^{-3}	8×10^{-3}	5–$6 \times 10^{-3\dagger}$
AdamW (β_1, β_2)	(0.8, 0.95)		
AdamW ϵ	10^{-10}		
AdamW weight decay	0		
Muon LR	8×10^{-3}	8×10^{-3}	6–$8 \times 10^{-3\dagger}$
Muon momentum	0.95		
Muon weight decay	0.2 (linear decay to 0)		
Muon NS steps	5		
LR schedule	Trapezoid (0% warmup, 50% cooldown)		
Min LR	0		
Gradient clipping	1.0		
Global batch size (tokens)	524K ($= 256 \times 2048$)		
Micro batch size	32	32	8–32†
Gradient checkpointing	✗	✗	model-dependent‡
`torch.compile`	✓*	✓*	✓*
Strategy	DDP		
Init strategy	scaled-zero + orthogonal		
Seed	42		

†At 770M, Attractor Model uses AdamW LR 5×10^{-3} and Muon LR 6×10^{-3}; Transformer/Parcae use 6×10^{-3} / 8×10^{-3}.
‡Gradient checkpointing enabled for 770M Attractor Model and Parcae; disabled for Transformer.
*`torch.compile` enabled for Transformer and Parcae; disabled for Attractor Model (implicit-gradient hooks are not compile-compatible).

Table 9. Parcae recurrence hyperparameters.

Hyperparameter	**140M**	**370M**	**770M**
Injection type	diagonal		
State init	like-init		
Mean recurrence	8	8	8
Mean backprop depth	4	4	4
Sampling scheme	poisson-truncated-full		
Iteration method	per-sequence		
Prelude norm	✓		

that trained it largely unnecessary at inference time. Interesting directions for future work is to further study the equilibrium internalization phenomenon and the differences between Attractor Models and finite-loop recurrence.