

Scaling Up Liquid-Resistance Liquid-Capacitance Networks for Efficient Sequence Modeling

Sequence modeling has recently been dominated by *linear* state-space models (SSMs), which can process very long sequences efficiently thanks to parallelization techniques. While *non-linear* SSMs are in principle more expressive, because they can capture richer input correlations through their internal states, their inherently sequential nature has made them much harder to scale. As a result, research has shifted away from non-linear models despite their potential advantages.

Recent work [3, 2] has made first steps toward scalable non-linear SSMs by adapting parallelization methods, but these solutions relied on approximations that either discarded useful information or limited scalability. In this work, we revisit the challenge of scaling non-linear models and introduce LrcSSM, a new bio-inspired recurrent architecture based on liquid-resistance liquid-capacitance (LRC) networks [1]. Our key idea is to redesign the state-transition dynamics so that the Jacobian matrix is inherently diagonal, while still being dependent on both the current state and the input. This simple but powerful change allows us to parallelize updates exactly, rather than approximately, while retaining the expressive benefits of non-linear recurrence.

From a theoretical perspective, LrcSSM offers both *computational efficiency* (matching the time and memory complexity of leading SSMs) and *formal stability guarantees*, something that other input-varying architectures such as Mamba or Liquid-S4 lack. Practically, this means the model can scale to long sequences with the same efficiency as linear models, but with the added ability to capture complex non-linear dynamics inspired by biological neurons.

We evaluate LrcSSM on standard long-sequence benchmarks from the UEA archive. Our results show that LrcSSM performs competitively across tasks, and in several cases surpasses strong SSM baselines such as LRU, S5, S6, and Mamba. In particular, on the EthanolConcentration benchmark, one of the most challenging due to intricate input correlations, LrcSSM achieves state-of-the-art performance.

In summary, this work makes the following contributions:

- We introduce LrcSSM, the first bio-inspired non-linear recurrent model that achieves the efficiency and scalability of modern SSMs.
- We demonstrate how inherent diagonalization enables exact parallelization of non-linear state updates.
- We provide formal stability guarantees alongside compute-optimal scaling behavior.
- We show empirically that LrcSSM matches or outperforms leading SSMs on long-sequence modeling tasks.

Overall, LrcSSM bridges a long-standing gap, it combines the *expressiveness* of non-linear dynamics with the *efficiency* of scalable SSMs, opening the door to a new class of biologically inspired, parallelizable architectures for sequence modeling.

- [1] Mónika Farsang, Sophie A Neubauer, and Radu Grosu. Liquid resistance liquid capacitance networks. In *The First Workshop on NeuroAI@ NeurIPS2024*, 2024.
- [2] Xavier Gonzalez, Andrew Warrington, Jimmy Smith, and Scott Linderman. Towards scalable and stable parallelization of nonlinear rnns. *Advances in Neural Information Processing Systems*, 37:5817–5849, 2024.
- [3] Yi Heng Lim, Qi Zhu, Joshua Selfridge, and Muhammad Firmansyah Kasim. Parallelizing non-linear sequential models over the sequence length. In *The Twelfth International Conference on Learning Representations*, 2024.