Dynamics and Representation Structure of Local Approximations to Gradient-Based Learning in Linear Recurrent Neural Networks

Ezekiel Williams, Alexandre Payeur, Guillaume Lajoie
✉ e.williams1@imperial.ac.uk

arXiv

Background

- Modern approaches to neuroscience use recurrent neural networks (RNNs) as models of brain circuits, and often hypothesize that biological learning should approximate gradients.
- RNN gradients are *non-local*: a parameter's derivative may depend on variables that are spatially or temporally far from the parameter. It is believed that biological learning is local.
- Computational neuroscientists have therefore formulated models for biological learning that locally approximate true gradients by removing non-local terms.

Research Problem and Objective

- Because they do not follow true gradients, there is limited understanding of when and how these local learning models will properly learn and, thus, of their strengths and weaknesses as models of biological learning.
- We study the local algorithm *random feedback online learning* (RFLO), in an analytically tractable RNN paradigm, to understand how its dynamics differ from backpropagation through time (BPTT) and truncated BPTT.

Theoretical Approach

1. We analyze a linear student-teacher model with hidden and output dimensions n and o.

Student-Teacher Model

$$h_{t+1} = Wh_t + Bx_t, \quad y_{t+1} = Ah_{t+1}, \qquad h^\star_{t+1} = W_\star h^\star_t + B_\star x_t, \quad y^\star_{t+1} = A_\star h^\star_{t+1}$$

Learning Task

$$x_t \sim \mathcal{N}(0,1), \quad \operatorname*{argmin}_\theta L(\theta),$$

$$L(\theta) = \frac{1}{2}\mathbb{E}\left(|y_T - y^\star_T|^2\right), \quad \theta = \{A, W, B\}$$

2. Following previous work[1], we assume a *mode-aligned* teacher and *data-aligned* student.

Assumption 1

$$W_0 = P\bar{W}_0P^\top, \quad A_0 = U\bar{A}_0P^\top, \quad B_0 = P\bar{B}_0V^\top$$

Assumption 2

$$W_\star = P_\star\bar{W}_\star P_\star^\top, \quad A_\star = U\bar{A}_\star P_\star^\top, \quad B_\star = P_\star\bar{B}_\star V_\star^\top$$

- Teacher model is mode aligned if it satisfies 2
- Student-teacher model is data aligned if it satisfies 1 and 2

3. With the data-aligned assumption, learning dynamics are fully described by those of A and B's signed singular values and W's eigenvalues: (a_i, b_i, w_i) $i \in \{0, \ldots, n\}$, with no coupling across all 3-dimensional modes. We drop the index and study an arbitrary mode.
4. We take $T \to \infty$, and the continuum limit of the learning dynamics, to obtain an ODE.
5. We use dynamical systems theory to derive learning solution manifolds (**A**), their stability (**B**), and asymptotic convergence rates.

ODE for RFLO*

$$\dot{a} = \frac{a_\star b b_\star}{1 - ww_\star} - \frac{ab^2}{1 - w^2}$$

$$\dot{b} = \frac{\hat{a}a_\star b_\star}{1 - \hat{w}w_\star} - \frac{\hat{a}ab}{1 - \hat{w}w}$$

$$\dot{w} = \frac{\hat{a}a_\star b b_\star w_\star}{(1 - \hat{w}w_\star)(1 - w_\star w)} - \frac{\hat{a}ab^2 w}{(1 - \hat{w}w)(1 - w^2)}$$

$\hat{a}$ and $\hat{w}$ depend on error feedback matrices

Without data-aligned assumptions, we separately prove that RFLO is constrained to learn perturbations of initial conditions that are restricted in rank (**D**):

Rank restriction

$$X_k = X_0 + \sum_{i=1}^{o} r_i q_i^\top, \; X \in \{W, B\}$$

r_i is row of error feedback matrix

*see paper for tBPTT and BPTT equations

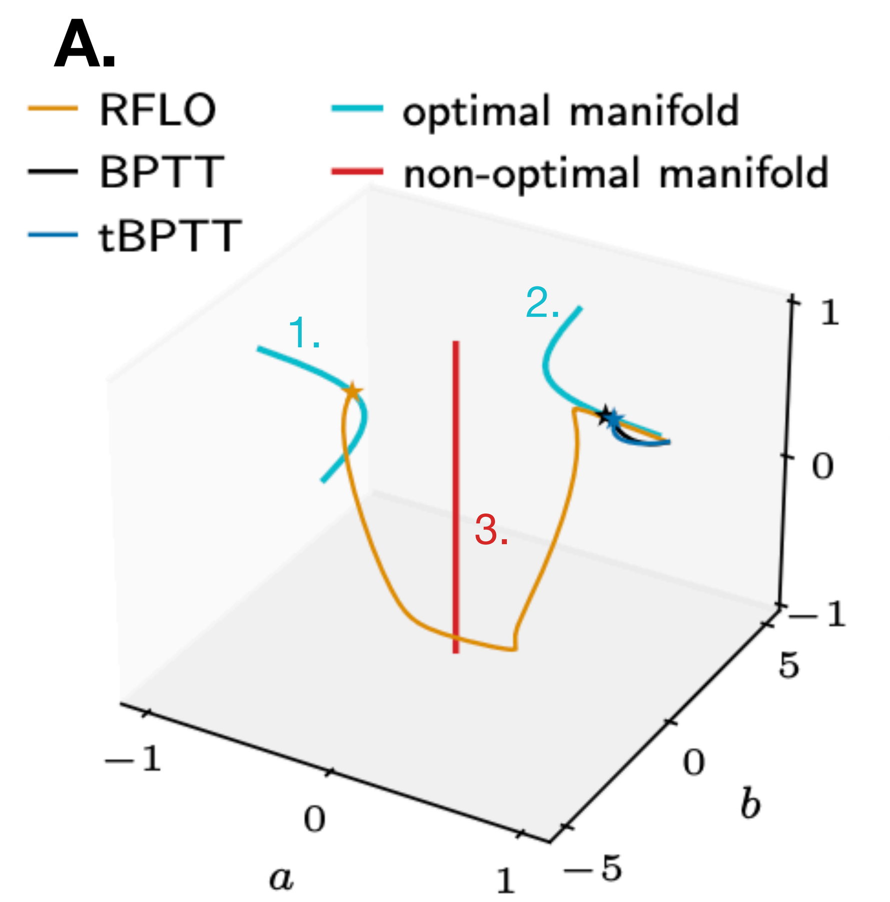

- BPTT & tBPTT: 1 & 2 stable; 3 unstable
- RFLO: 1 stable; 2 partially unstable

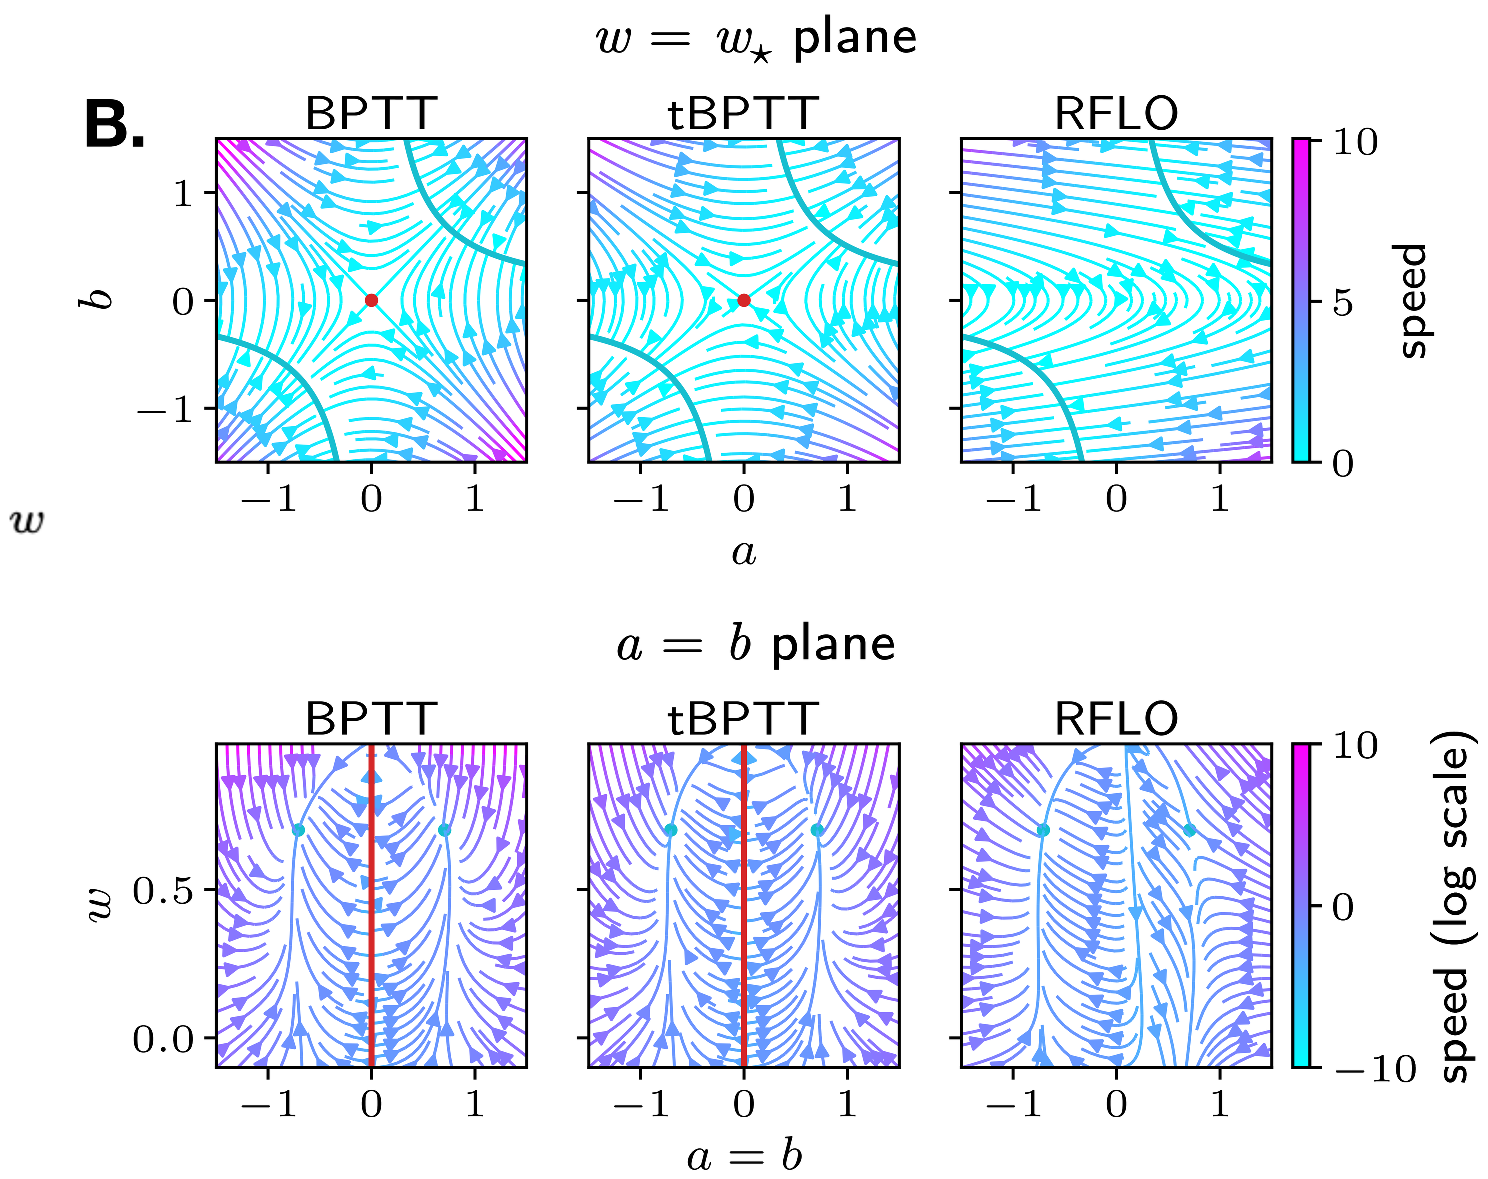

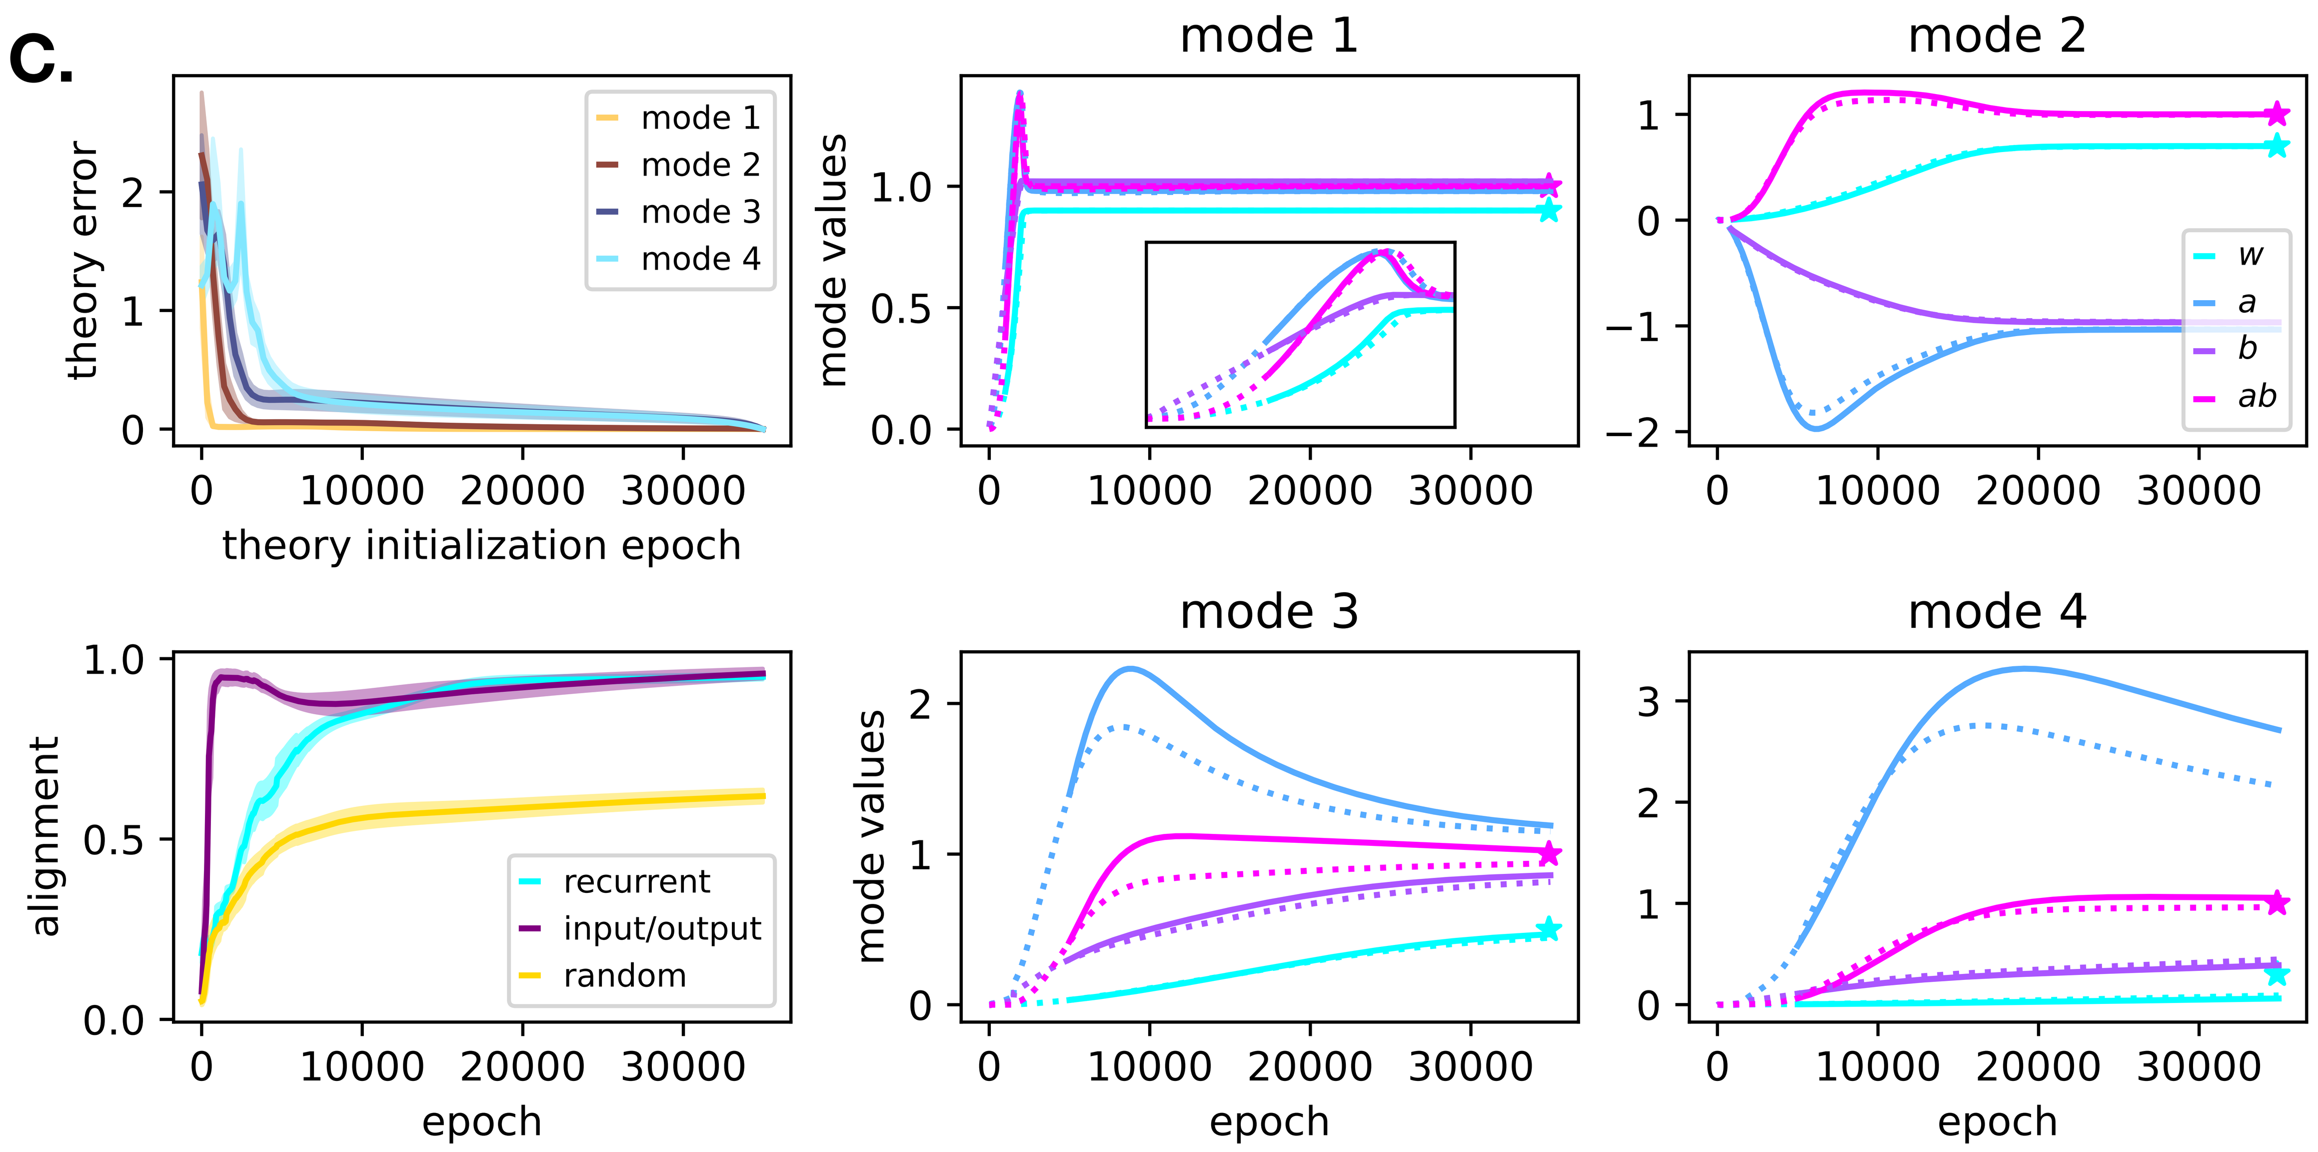

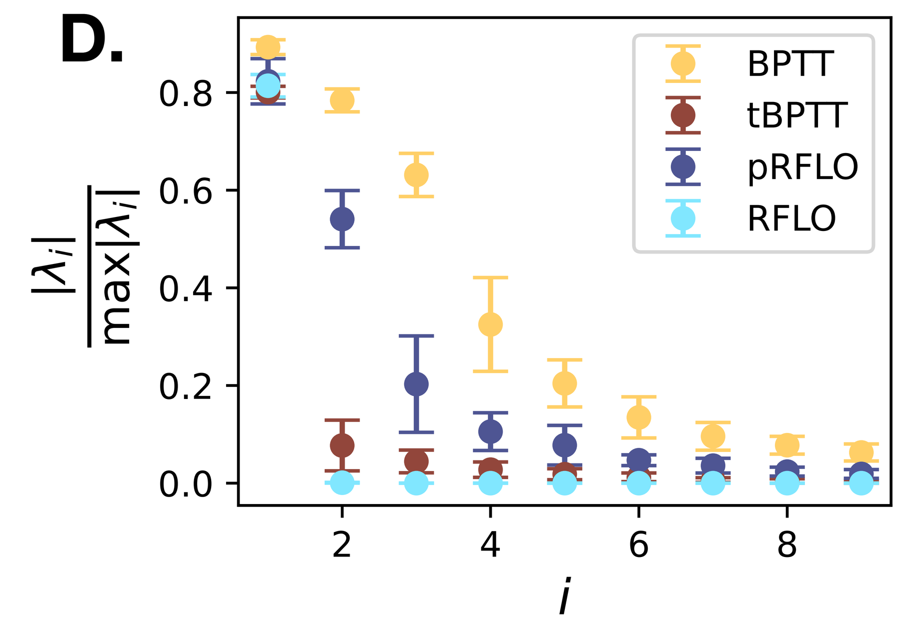

A. Manifold structure of one parameter mode triplet (a, b, w), with example dynamics

B. Vector fields on 2-dimensional slices through parameter space for one mode, showing manifold stability

C. Numerical validation: theory provides insights for mode-aligned teacher even when student is not data-aligned

D. Local learning algorithms learn low-rank representations

Conclusions and Future Directions

- Similarity between tBPTT and BPTT learned manifolds suggests disentangling dynamic modes, not just length of task timescales, might determine difficulty of temporal credit assignment problem.
- In our model, RFLO has fewer stable optima and slower asymptotic convergence rates (see paper). We speculate that similar phenomena may impact local learning in more sophisticated models.
- Future research needs to go beyond the data-alignment assumption.
- Low-rank constraint for RFLO (and, potentially, tBPTT and e-prop) may prohibit learning in certain cases. For example, when there are more task timescales than output dimensions.

1. Proca, A. M. et al. *Learning dynamics in linear recurrent neural networks.* ICML 2025.