Novelty Assessment Report

Paper: Temporal superposition and feature geometry of RNNs under memory demands

PDF URL: https://openreview.net/pdf?id=7cMzTpbJHC

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

Understanding how populations of neurons represent information is a central challenge across machine learning and neuroscience. Recent work in both fields has begun to characterize the representational geometry and functionality underlying complex distributed activity. For example, artificial neural networks trained on data with more features than neurons compress data by representing features non-orthogonally in so-called superposition. However, the effect of time (or memory), an additional capacity-constraining pressure, on underlying representational geometry in recurrent models is not well understood. Here, we study how memory demands affect representational geometry in recurrent neural networks (RNNs), introducing the concept of temporal superposition. We develop a theoretical framework in RNNs with linear recurrence trained on a delayed serial recall task to better understand how properties of the data, task demands, and network dimensionality lead to different representational strategies, and show that these insights generalize to nonlinear RNNs. Through this, we identify an effectively linear, dense regime and a sparse regime where RNNs utilize an interference-free space, characterized by a phase transition in the angular distribution of features and decrease in spectral radius. Finally, we analyze the interaction of spatial and temporal superposition to observe how RNNs mediate different representational tradeoffs. Overall, our work offers a mechanistic, geometric explanation of representational strategies RNNs learn, how they depend on capacity and task demands, and why.

Core Task Landscape

This paper addresses: **Representational Geometry in Recurrent Neural Networks under Memory Constraints**

A total of **50 papers** were analyzed and organized into a taxonomy with **12 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Representational Geometry and Memory Organization**
- **Learning Dynamics and Representational Development**
- **Architecture Design and Memory Mechanisms**
- **Task-Specific Dynamics and Applications**
- **Theoretical Foundations and Computational Principles**

Complete Taxonomy Tree

- Representational Geometry in Recurrent Neural Networks under Memory Constraints Survey Taxonomy
- Representational Geometry and Memory Organization
  - Geometric Properties of Memory Representations ★ (6 papers)
  - [0] Temporal superposition and feature geometry of RNNs under memory demands (Anon et al., 2026) View paper
  - [1] Measuring and controlling solution degeneracy across task-trained recurrent neural networks (Huang Ann, 2025) View paper
  - [2] Geometry of naturalistic object representations in recurrent neural network models of working memory (Pouya Bashivan, 2024) View paper
  - [7] Geometry of neural computation unifies working memory and planning (Daniel B. Ehrlich, 2022) View paper
  - [23] Geometry and dynamics of representations in a precisely balanced memory network related to olfactory cortex. (Claire Meissner-Bernard, 2025) View paper
  - [24] Recurrent neural network models for working memory of continuous variables: activity manifolds, connectivity patterns, and dynamic codes (Christopher J. Cueva, 2021) View paper
  - Memory Capacity and Information Storage (6 papers)
  - [15] Capacity and trainability in recurrent neural networks (Jasmine Collins, 2016) View paper
  - [18] Memory and information processing in recurrent neural networks (Alireza Goudarzi, 2016) View paper
  - [22] Memory capacity of recurrent neural networks with matrix representation (Renanse, 2023) View paper
  - [30] Distributed Sequence Memory of Multidimensional Inputs in Recurrent Networks (Adam S. Charles, 2022) View paper
  - [36] Optimal short-term memory before the edge of chaos in driven random recurrent networks (Haruna, 2022) View paper
  - [42] The difference between memory and prediction in linear recurrent networks (Marzen, 2022) View paper
  - Representational Degeneracy and Solution Variability (2 papers)
  - [8] Representational learning by optimization of neural manifolds in an olfactory memory network (Rainer Friedrich, 2024) View paper
  - [16] Sudden restructuring of memory representations in recurrent neural networks with repeated stimulus presentations. (Howlett, 2025) View paper
- Learning Dynamics and Representational Development
  - Learning Trajectories and Phase Transitions (3 papers)
  - [3] Learning dynamics and the geometry of neural dynamics in recurrent neural controllers (A Huang, 2024) View paper
  - [5] Understanding and controlling the geometry of memory organization in RNNs (Haputhanthri, 2025) View paper

- [26] Learning Dynamics and Geometry in Recurrent Neural Controllers (A Huang, 2024) View paper
  - Optimization and Constraint-Based Learning (3 papers)
  - [20] Geometry Constrained Progressive Learning for Lstm-Based Speech Enhancement (Tang Xin, 2020) View paper
  - [21] Learning and development in neural networks: The importance of starting small (ELMAN, 1993) View paper
  - [35] Fading memory as inductive bias in residual recurrent networks (Igor Dubinin, 2023) View paper
- Architecture Design and Memory Mechanisms
  - Specialized Memory Structures (6 papers)
  - [6] Investigating recurrent neural network memory structures using neuro-evolution (Ororbia, 2019) View paper
  - [14] A theory of sequence indexing and working memory in recurrent neural networks (Frady, 2018) View paper
  - [25] Hebbian Memory-Augmented Recurrent Networks: Engram Neurons in Deep Learning (Szelogowski, 2025) View paper
  - [37] Learning context-free grammars: Capabilities and limitations of a recurrent neural network with an external stack memory (Sreerupa Das, 1992) View paper
  - [41] Rapid memory encoding in a recurrent network model with behavioral time scale synaptic plasticity (Pan Ye, 2023) View paper
  - [50] Learning to Control Rapidly Changing Synaptic Connections: An Alternative Type of Memory in Sequence Processing Artificial Neural Networks (Irie, 2022) View paper
  - Standard Architecture Analysis (4 papers)
  - [13] Recurrent neural networks and their memory behavior: a survey (Yuan-Hang Su, 2022) View paper
  - [27] Recurrent neural networks with finite memory length (Ding-Kun Long, 2019) View paper
  - [33] How embedded memory in recurrent neural network architectures helps learning long-term temporal dependencies (Tsung-Nan Lin, 1998) View paper
  - [38] A formal hierarchy of RNN architectures (William Merrill, 2020) View paper
  - Gated and Attention-Based Mechanisms (4 papers)
  - [40] Attention-based Memory Selection Recurrent Network for Language Modeling (Liu, 2022) View paper
  - [47] Memory-Gated Recurrent Networks (Dai Min, 2025) View paper
  - [48] Self-Gated Memory Recurrent Network for Efficient Scalable HDR Deghosting (Prabhakar, 2021) View paper
  - [49] Scalable Online Recurrent Learning Using Columnar Neural Networks (Javed, 2021) View paper
- Task-Specific Dynamics and Applications
  - Working Memory Tasks (5 papers)
  - [4] A recurrent neural network model of prefrontal brain activity during a working memory task (Emilia P. Piwek, 2023) View paper
  - [12] Engineering recurrent neural networks from task-relevant manifolds and dynamics (Eli Pollock, 2020) View paper
  - [19] Oscillatory control over representational geometry of sequence working memory in macaque frontal cortex (Fang Wen, 2025) View paper
  - [29] Trained recurrent neural networks develop phase-locked limit cycles in a working memory task. (Matthijs Pals, 2024) View paper
  - [46] Phase remembers: trained RNNs develop phase-locked limit cycles in a working memory task (Matthijs Pals, 2023) View paper
  - Navigation and Spatial Reasoning (4 papers)
  - [9] Learning spatial common sense with geometry-aware recurrent networks (Hsiao-Yu Fish Tung, 2019) View paper
  - [28] 3D View Prediction Models of the Dorsal Visual Stream (Gabriel Sarch, 2023) View paper
  - [31] eLife assessment: Remapping in a recurrent neural network model of navigation and context inference (Srdjan Ostojic, 2023) View paper
  - [34] Author Response: Remapping in a recurrent neural network model of navigation and context inference (Isabel I.C. Low, 2023) View paper
  - Sequence Processing and Linguistic Tasks (5 papers)
  - [11] Distributed representations, simple recurrent networks, and grammatical structure (Jeffrey L. Elman, 1991) View paper
  - [17] Short-term memory for serial order: a recurrent neural network model. (Botvinick, 2006) View paper
  - [39] Multi-view text classification through integrated RNN autoencoder learning of word, sentence, emotion and paragraph representations (Yitao Ding, 2025) View paper
  - [43] Empirical Analysis of Limits for Memory Distance in Recurrent Neural Networks (Illium, 2022) View paper
  - [44] State-Regularized Recurrent Neural Networks to Extract Automata and Explain Predictions (Wang Cheng, 2022) View paper
- Theoretical Foundations and Computational Principles
  - Dynamical Systems and Computational Theory (3 papers)
  - [10] Structural determinants of soft memory in recurrent biological networks (Maria Sol Vidal-Saez, 2025) View paper
  - [32] Understanding computation through low-dimensional dynamics with recurrent neural networks (Pollock, 2022) View paper
  - [45] Recurrent Network Models Of Sequence Generation And Memory (Rajan, 2022) View paper

Narrative

Core task: representational geometry in recurrent neural networks under memory constraints. This field examines how RNNs organize and structure internal representations when faced with limited memory resources, spanning questions about geometric properties of neural codes, learning dynamics that shape these representations, architectural innovations for memory management, task-specific adaptations, and underlying theoretical principles. The taxonomy reflects five major branches: Representational Geometry and Memory Organization explores how memory states are geometrically arranged and what structural properties emerge (e.g., Naturalistic Object Geometry[2], Balanced Memory Geometry[23]); Learning Dynamics and Representational Development investigates how training shapes these geometries over time (e.g., Learning Dynamics Geometry[3]); Architecture Design and Memory Mechanisms focuses on novel gating structures and memory augmentation strategies (e.g., Memory Gated Recurrent[47], Hebbian Memory Augmented[25]); Task-Specific Dynamics examines how different cognitive demands—navigation, sequence processing, working memory—drive distinct representational solutions (e.g., Prefrontal Working Memory[4], Unified Working Memory[7]); and Theoretical Foundations provides capacity analyses and computational principles (e.g., Matrix Representation Capacity[22], Capacity and Trainability[15]).

Several active lines of work reveal key trade-offs and open questions. One cluster investigates how networks balance memory capacity with geometric structure: some studies emphasize controlling solution degeneracy and maintaining balanced representations (Solution Degeneracy Control[1], Memory Geometry Control[5]), while others explore how oscillatory dynamics or phase-based codes can expand representational capacity (Oscillatory Representational Geometry[19]). Another thread examines the interplay between task demands and emergent geometry, asking whether memory constraints force networks into low-dimensional manifolds or enable richer, task-relevant structures (Task Relevant Manifolds[12], Activity Manifolds Connectivity[24]). Temporal Superposition RNNs[0] sits within the Geometric Properties of Memory Representations cluster, closely aligned with works like Balanced Memory Geometry[23] that study how networks organize overlapping temporal information in constrained state spaces. Compared to Memory Geometry Control[5], which focuses on explicit geometric constraints during training, Temporal Superposition RNNs[0] emphasizes the natural emergence of superposed codes under memory pressure, offering a complementary perspective on how recurrent architectures handle multiple simultaneous memory demands.

Related Works in Same Category

The following **5 sibling papers** share the same taxonomy leaf node with the original paper:

1. Measuring and controlling solution degeneracy across task-trained recurrent neural networks

Authors: Huang Ann, Martinelli, Flavio, Rajan, Kanaka | **Year/Venue**: 2025 | **URL**: View paper

Abstract

â¦ features, their representational geometry can vary widely. â¦ the same constraints on representational degeneracy. â¦ only memory-intensive task: our method estimates a memory â¦

Relationship Analysis

Both papers belong to the same taxonomy category investigating geometric properties of memory representations in RNNs. They share overlapping interests in how RNNs organize representations under memory constraints, with both examining manifold structure and spatial organization of learned features. However, the original paper focuses specifically on temporal superposition and feature geometry under varying memory demands through controlled delay tasks, while the candidate paper addresses solution degeneracy across different task-trained RNNs and methods for measuring and controlling representational variability.

---

2. Geometry of naturalistic object representations in recurrent neural network models of working memory

Authors: Pouya Bashivan, Takuya Ito, Xiaoxuan Lei | **Year/Venue**: 2024 • Neural Information Processing Systems | **URL**: View paper

Abstract

Working memory is a central cognitive ability crucial for intelligent decision-making. Recent experimental and computational work studying working memory has primarily used categorical (i.e., one-hot) inputs, rather than ecologically relevant, multidimensional naturalistic ones. Moreover, studies have primarily investigated working memory during single or few cognitive tasks. As a result, an understanding of how naturalistic object information is maintained in working memory in neural networks i...

Relationship Analysis

Both papers belong to the same taxonomy category investigating geometric properties of memory representations in RNNs, sharing a focus on how neural networks organize information in their latent spaces. While the original paper examines temporal superposition and feature geometry under controlled memory constraints using synthetic delay tasks, the candidate paper studies naturalistic object representations in working memory across multiple N-back tasks with real-world stimuli. The key difference is that the original paper develops theoretical frameworks for understanding temporal compression mechanisms in minimal settings, whereas the candidate paper empirically analyzes how task-relevant and irrelevant naturalistic features are maintained across different cognitive tasks in sensory-cognitive architectures.

---

3. Geometry of neural computation unifies working memory and planning

Authors: Daniel B. Ehrlich, John D. Murray | **Year/Venue**: 2022 | **URL**: View paper

Abstract

Real-world tasks require coordination of working memory, decision-making, and planning, yet these cognitive functions have disproportionately been studied as independent modular processes in the brain. Here, we propose that contingency representations, defined as mappings for how future behaviors depend on upcoming events, can unify working memory and planning computations. We designed a task capable of disambiguating distinct types of representations. In task-optimized recurrent neural networks...

Relationship Analysis

Both papers belong to the Geometric Properties of Memory Representations category, investigating how recurrent neural networks organize and represent information in activation space under memory constraints. While the original paper focuses on temporal superposition and how memory demands affect feature geometry through controlled delay tasks, the candidate paper examines contingency representations that unify working memory and planning by mapping future behaviors to upcoming events. The key difference is that the original paper characterizes geometric strategies (spiral arrangements, interference-free spaces) for temporal feature compression, whereas the candidate paper proposes a unified computational framework for working memory and planning through contingency mappings.

---

4. Geometry and dynamics of representations in a precisely balanced memory network related to olfactory cortex.

Authors: Claire Meissner-Bernard, Friedemann Zenke, Rainer W Friedrich | **Year/Venue**: 2025 • eLife | **URL**: View paper

Abstract

N/A

Relationship Analysis

Both papers belong to the same taxonomy category investigating geometric properties of memory representations in RNNs, specifically examining manifold structure and spatial organization of neural representations. The candidate paper focuses on geometry and dynamics in a precisely balanced memory network related to olfactory cortex, likely emphasizing biological neural circuits and their representational properties. The original paper differs by introducing temporal superposition as a novel concept, analyzing how memory demands affect representational geometry through theoretical frameworks on delayed recall tasks, and identifying phase transitions between dense and sparse regimes in artificial RNNs rather than biological networks.

---

5. Recurrent neural network models for working memory of continuous variables: activity manifolds, connectivity patterns, and dynamic codes

Authors: Christopher J. Cueva, Ardalan, Adel, Adel Ardalan, Tsodyks Misha, et al. (8 authors total) | **Year/Venue**: 2021 | **URL**: View paper

Abstract

Many daily activities and psychophysical experiments involve keeping multiple items in working memory. When items take continuous values (e.g., orientation, contrast, length, loudness) they must be stored in a continuous structure of appropriate dimensions. We investigate how this structure is represented in neural circuits by training recurrent networks to report two previously shown stimulus orientations. We find the activity manifold for the two orientations resembles a Clifford torus. Althou...

Relationship Analysis

Both papers belong to the same taxonomy category investigating geometric properties of memory representations in RNNs, specifically examining manifold structure and spatial organization of neural activity. They share overlapping interests in how RNNs represent continuous variables in memory through geometric analysis of hidden state representations, using dimensionality reduction and manifold characterization techniques. However, the original paper focuses on temporal superposition and how memory demands over time affect feature geometry (introducing concepts like temporal vs. spatial superposition and interference-free spaces), while the candidate paper examines the specific geometric structure (Clifford torus) that emerges when RNNs store multiple continuous orientations simultaneously, emphasizing connectivity patterns and dynamic coding schemes rather than temporal capacity constraints.

Contributions Analysis

Overall novelty summary. The paper introduces temporal superposition as a framework for understanding how recurrent neural networks represent multiple features under memory constraints, focusing on delayed serial recall tasks. It resides in the 'Geometric Properties of Memory Representations' leaf, which contains six papers examining manifold structure and spatial organization of memory codes. This leaf sits within the broader 'Representational Geometry and Memory Organization' branch, indicating a moderately populated research direction. The sibling papers address related geometric questions—balanced memory structures, naturalistic object geometry, and manifold connectivity—suggesting the paper enters an active but not overcrowded subfield where geometric analysis of memory is an established concern.

The taxonomy reveals neighboring research directions that contextualize this work. The adjacent 'Memory Capacity and Information Storage' leaf (six papers) focuses on theoretical bounds rather than geometric structure, while 'Learning Dynamics and Representational Development' (six papers across two leaves) examines how representations evolve during training rather than their static properties. The 'Task-Specific Dynamics' branch includes a 'Working Memory Tasks' leaf (five papers) studying similar cognitive paradigms but emphasizing task performance over geometric principles. The paper's focus on geometric organization under temporal constraints bridges these areas, connecting capacity theory with representational structure in a way that distinguishes it from purely capacity-focused or purely task-driven analyses.

Among seventeen candidates examined, three contributions show evidence of prior overlap. The temporal superposition concept (ten candidates examined, one refutable) appears to have some precedent in the limited search scope, though nine candidates did not clearly refute it. The theoretical framework with loss decomposition (five candidates, one refutable) and the identification of interference-free regimes (two candidates, one refutable) each face one potentially overlapping prior work among the small candidate pools examined. These statistics suggest that while the core ideas have some grounding in existing literature, the specific formulation and integration may offer incremental advances. The limited search scope—seventeen total candidates—means these assessments reflect top semantic matches rather than exhaustive coverage.

Based on the constrained literature search, the work appears to synthesize existing geometric and capacity concerns into a unified temporal framework. The taxonomy position indicates a moderately active research area with clear boundaries separating geometric analysis from learning dynamics and architectural design. The contribution-level statistics suggest partial novelty: each major claim encounters at least one potentially overlapping candidate among the limited pool examined, but substantial portions of the candidate sets do not clearly refute the contributions. This pattern is consistent with incremental theoretical refinement rather than a foundational shift, though the restricted search scope limits definitive conclusions about the work's broader originality.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Concept of temporal superposition in RNNs

Description: The authors introduce temporal superposition as a novel form of representational compression in recurrent neural networks that arises from memory demands. Unlike spatial superposition (compressing more input features than neurons), temporal superposition occurs when features must be maintained over time longer than the hidden state dimensionality allows, forcing the network to represent features non-orthogonally across temporal positions.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Recurrent neural networks for edge intelligence: a survey

URL: View paper

Brief Assessment

Edge Intelligence Survey[61] focuses on RNN applications for edge computing and compression techniques. It does not address temporal superposition as a form of representational compression arising from memory demands in recurrent networks.

---

2. Comparative study of state-based neural networks for virtual analog audio effects modeling

URL: View paper

Brief Assessment

Virtual Analog Modeling[58] focuses on applying recurrent neural networks (RNNs, LSTMs, SSMs) to audio effects modeling tasks, not on theoretical analysis of representational geometry or temporal superposition concepts in RNNs under memory constraints.

---

3. A temporal convolutional recurrent autoencoder based framework for compressing time series data

URL: View paper

Brief Assessment

The candidate paper (Temporal Convolutional Autoencoder[57]) focuses on time series data compression using temporal convolutional networks combined with recurrent neural networks for encoding/decoding. It does not address representational geometry, memory-induced compression strategies, or the concept of temporal superposition as defined in the original work.

---

4. Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task

URL: View paper

Brief Assessment

Physical Inference Patterns[62] focuses on comparing RNN behavior to primates in a physical inference task using dynamic latent variable tracking, not on temporal superposition as a form of representational compression arising from memory demands in RNNs.

---

5. Temporal superimposed crossover module for effective continuous sign language

URL: View paper

Brief Assessment

Temporal Superimposed Crossover[60] addresses continuous sign language recognition using temporal superposition operations on video frames, not representational compression in recurrent neural networks under memory demands.

---

6. Fostering Event-Predictive Encodings in Recurrent Neural Networks

URL: View paper

Brief Assessment

Event Predictive Encodings[65] focuses on event segmentation and prediction in hierarchical architectures using surprise signals, not on representational compression strategies in RNNs under memory constraints. The candidate addresses temporal event boundaries rather than temporal superposition of features in recurrent networks.

---

7. Bbs-rnn: Block-based Structure Compression With Admm For Rnn On Temporal Sequence Applications.

URL: View paper

Brief Assessment

Block Based Compression[64] focuses on block-based structure compression techniques for RNNs using ADMM optimization methods, which is a hardware/compression approach rather than a theoretical framework about representational geometry and temporal superposition in neural networks.

---

8. An extended echo state network using Volterra filtering and principal component analysis

URL: View paper

Brief Assessment

Volterra Echo State[63] focuses on echo state networks using Volterra filtering and PCA for nonlinear system modeling. The candidate's brief mention of 'temporal superposition' appears in a different technical context (signal processing/compression) rather than the original paper's framework of representational geometry and memory-induced feature compression in RNNs.

---

9. THE COMPUTATIONAL ROLE OF COMPLEX REPRESENTATIONS IN RNNS

URL: View paper

Prior Art Analysis

Complex Representations Role[55] demonstrates that the concept of temporal superposition in RNNs was previously introduced and analyzed. The candidate paper explicitly states it investigates 'internal representations of rnns from the perspective of superposition' and develops 'a theoretical framework that predicts and explains the geometric arrangement of features in the hidden states of both linear and non-linear rnns.' The candidate's abstract describes the same core concept: representing features across time in a capacity-constrained hidden state, leading to non-orthogonal feature representations. Both papers study how memory demands force RNNs to compress temporal features into lower-dimensional spaces.

Evidence

Evidence 1 - **Rationale**: Both papers explicitly study RNN internal representations through the lens of superposition, developing theoretical frameworks to explain feature geometry under memory constraints. - **Original**: here, we study how memory demands affect representational geometry in recurrent neural networks (rnns), introducing the concept of temporal superposition. - **Candidate**: w e investigate the internal representations of rnns from the perspective of superposition. beginning with an exact analytical decomposition of the expected loss on a recall task, we develop a theoretical framework that predicts and explains the geometric arrangement of features in the hidden states...

Evidence 2 - **Rationale**: The candidate acknowledges that superposition research has not been extended to RNNs, but the original paper explicitly introduces temporal superposition as a novel concept for RNNs, suggesting the candidate's work predates this contribution. - **Original**: temporal superposition. in addition to spatial superposition, we claim that rnns exhibit another form of superposition due to the axis of time and that this phenomenon is fundamentally different from the feedforward case. in particular, in addition to having some spatial component, we can think of e... - **Candidate**: the superposition hypothesis posits that this occurs because models learn to represent more features than they have neurons. most research on superposition focuses on feed-forward models and the implications of this hypothesis have not been extended to more complex architectures, such as rnns.

---

10. HT-STNet: a hierarchical Tucker decomposition and spatio-temporal LSTM network for accurate and efficient shared mobility demand forecasting on sparse data

URL: View paper

Brief Assessment

Hierarchical Tucker Decomposition[59] focuses on tensor decomposition methods for spatio-temporal forecasting in mobility demand prediction, not on representational compression or memory-induced superposition in recurrent neural networks.

Contribution 2: Theoretical framework with loss decomposition

Description: The authors derive an analytical expression for the loss on a k-delay task that decomposes into four interpretable terms: task benefit, mean correction, projection interference cost, and composition interference. This decomposition explains the geometric strategies employed by RNNs and how data properties and network dimensionality interact with memory demands.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Resource-Efficient Acoustic Full-Waveform Inversion via Dual-Branch Physics-Informed RNN with Scale Decomposition

URL: View paper

Brief Assessment

Dual Branch Physics[52] focuses on full-waveform seismic inversion using dual-branch physics-informed RNNs with scale decomposition for resource efficiency. It does not address theoretical frameworks for RNN memory strategies or loss decomposition into interpretable terms related to memory demands and feature geometry.

---

2. PredRNN: A Recurrent Neural Network for Spatiotemporal Predictive Learning

URL: View paper

Brief Assessment

PredRNN[51] focuses on spatiotemporal video prediction with memory flow architectures, not on theoretical loss decomposition for RNN memory strategies or analytical frameworks explaining geometric strategies under memory demands.

---

3. Two-shot learning of continuous interpolation using a conceptor-aided recurrent autoencoder

URL: View paper

Brief Assessment

Conceptor Recurrent Autoencoder[54] focuses on two-shot learning for temporal pattern interpolation using conceptor-based regularization, not on analytical loss decomposition for understanding RNN memory strategies or geometric feature representations under memory demands.

4. THE COMPUTATIONAL ROLE OF COMPLEX REPRESENTATIONS IN RNNS

URL: View paper

Prior Art Analysis

Complex Representations Role[55] presents an analytical loss decomposition framework prior to the original paper. The candidate explicitly states it begins 'with an exact analytical decomposition of the expected loss on a recall task' to develop its theoretical framework. Both papers derive expressions that decompose the loss into interpretable terms to explain RNN geometric strategies. The candidate's abstract directly describes the same methodological approach of using loss decomposition to understand feature geometry.

Evidence

Evidence 1 - **Rationale**: Both papers introduce and distinguish between projection and composition interference as key components of their theoretical frameworks. - **Original**: we distinguish two forms of interference - projection and composition interference - and show how they impact behavior. - **Candidate**: w e introduce a key distinction between projection interference and composition interference and use this to explain why , under high sparsity , the optimal strategy for non-linear models is to pack most of the task-relevant features into just half of the activation space.

5. Predicting Wave Dynamics using Deep Learning with Multistep Integration Inspired Attention and Physics-Based Loss Decomposition

URL: View paper

Brief Assessment

Wave Dynamics Prediction[53] focuses on decomposing loss for wave propagation PDEs into dissipation and dispersion components, not on analyzing RNN memory strategies or geometric feature representations in recurrent networks.

Contribution 3: Identification of interference-free space and phase transition

Description: The authors identify that RNNs with nonlinear readouts can exploit an interference-free space (the half-space opposite the readout direction) to pack intermediate feature directions without projection interference. They characterize a phase transition between dense and sparse regimes marked by changes in angular distribution of features and spectral radius, with nonlinear RNNs implementing sharp forgetting by fully exploiting this space.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Predicting change points in multivariate time series data

URL: View paper

Brief Assessment

Change Point Prediction[56] focuses on detecting change points in multivariate time series (e.g., seizure prediction from EEG data) rather than analyzing RNN feature geometry or interference-free spaces in neural network representations.

2. THE COMPUTATIONAL ROLE OF COMPLEX REPRESENTATIONS IN RNNS

URL: View paper

Prior Art Analysis

Complex Representations Role[55] identifies the interference-free space concept and its exploitation by RNNs before the original paper. The candidate's abstract explicitly states it explains 'why, under high sparsity, the optimal strategy for non-linear models is to pack most of the task-relevant features into just half of the activation space.' This directly corresponds to the original paper's interference-free space concept. Both papers analyze how nonlinear RNNs exploit this half-space to minimize interference and describe the geometric strategies that emerge.

Evidence

Evidence 1 - **Rationale**: Both papers identify that nonlinear RNNs optimally pack features into half the activation space (the interference-free space) under high sparsity conditions. - **Original**: in rnns with relu nonlinearities, we identify the existence of an interference-free space into which many feature directions can be tightly packed. we identify a phase transition in the geometry between the dense and sparse regimes. - **Candidate**: w e introduce a key distinction between projection interference and composition interference and use this to explain why , under high sparsity , the optimal strategy for non-linear models is to pack most of the task-relevant features into just half of the activation space.

Evidence 2 - **Rationale**: Both papers derive loss approximations under high sparsity to explain the geometric strategies, including the interference-free space exploitation. - **Original**: the nonlinear setting makes deriving analytic solutions to the loss more challenging. we therefore approximate the expectation of the loss in the limit of high temporal sparsity (appendix e.1), yielding e[l] ≈ pν tx t=k+1 (relu(w⊤ y ws=k) - 1)2 | {z } task benefit + tx t=1 t-1x s̸=k relu(w⊤ y ws)2 |... - **Candidate**: beginning with an exact analytical decomposition of the expected loss on a recall task, we develop a theoretical framework that predicts and explains the geometric arrangement of features in the hidden states of both linear and non-linear rnns.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Temporal superposition and feature geometry of RNNs under memory demands View paper
- [1] Measuring and controlling solution degeneracy across task-trained recurrent neural networks View paper
- [2] Geometry of naturalistic object representations in recurrent neural network models of working memory View paper
- [3] Learning dynamics and the geometry of neural dynamics in recurrent neural controllers View paper
- [4] A recurrent neural network model of prefrontal brain activity during a working memory task View paper
- [5] Understanding and controlling the geometry of memory organization in RNNs View paper
- [6] Investigating recurrent neural network memory structures using neuro-evolution View paper
- [7] Geometry of neural computation unifies working memory and planning View paper
- [8] Representational learning by optimization of neural manifolds in an olfactory memory network View paper
- [9] Learning spatial common sense with geometry-aware recurrent networks View paper
- [10] Structural determinants of soft memory in recurrent biological networks View paper

- [11] Distributed representations, simple recurrent networks, and grammatical structure View paper
- [12] Engineering recurrent neural networks from task-relevant manifolds and dynamics View paper
- [13] Recurrent neural networks and their memory behavior: a survey View paper
- [14] A theory of sequence indexing and working memory in recurrent neural networks View paper
- [15] Capacity and trainability in recurrent neural networks View paper
- [16] Sudden restructuring of memory representations in recurrent neural networks with repeated stimulus presentations. View paper
- [17] Short-term memory for serial order: a recurrent neural network model. View paper
- [18] Memory and information processing in recurrent neural networks View paper
- [19] Oscillatory control over representational geometry of sequence working memory in macaque frontal cortex View paper
- [20] Geometry Constrained Progressive Learning for Lstm-Based Speech Enhancement View paper
- [21] Learning and development in neural networks: The importance of starting small View paper
- [22] Memory capacity of recurrent neural networks with matrix representation View paper
- [23] Geometry and dynamics of representations in a precisely balanced memory network related to olfactory cortex. View paper
- [24] Recurrent neural network models for working memory of continuous variables: activity manifolds, connectivity patterns, and dynamic codes View paper
- [25] Hebbian Memory-Augmented Recurrent Networks: Engram Neurons in Deep Learning View paper
- [26] Learning Dynamics and Geometry in Recurrent Neural Controllers View paper
- [27] Recurrent neural networks with finite memory length View paper
- [28] 3D View Prediction Models of the Dorsal Visual Stream View paper
- [29] Trained recurrent neural networks develop phase-locked limit cycles in a working memory task. View paper
- [30] Distributed Sequence Memory of Multidimensional Inputs in Recurrent Networks View paper
- [31] eLife assessment: Remapping in a recurrent neural network model of navigation and context inference View paper
- [32] Understanding computation through low-dimensional dynamics with recurrent neural networks View paper
- [33] How embedded memory in recurrent neural network architectures helps learning long-term temporal dependencies View paper
- [34] Author Response: Remapping in a recurrent neural network model of navigation and context inference View paper
- [35] Fading memory as inductive bias in residual recurrent networks View paper
- [36] Optimal short-term memory before the edge of chaos in driven random recurrent networks View paper
- [37] Learning context-free grammars: Capabilities and limitations of a recurrent neural network with an external stack memory View paper
- [38] A formal hierarchy of RNN architectures View paper
- [39] Multi-view text classification through integrated RNN autoencoder learning of word, sentence, emotion and paragraph representations View paper
- [40] Attention-based Memory Selection Recurrent Network for Language Modeling View paper
- [41] Rapid memory encoding in a recurrent network model with behavioral time scale synaptic plasticity View paper
- [42] The difference between memory and prediction in linear recurrent networks View paper
- [43] Empirical Analysis of Limits for Memory Distance in Recurrent Neural Networks View paper
- [44] State-Regularized Recurrent Neural Networks to Extract Automata and Explain Predictions View paper
- [45] Recurrent Network Models Of Sequence Generation And Memory View paper
- [46] Phase remembers: trained RNNs develop phase-locked limit cycles in a working memory task View paper
- [47] Memory-Gated Recurrent Networks View paper
- [48] Self-Gated Memory Recurrent Network for Efficient Scalable HDR Deghosting View paper
- [49] Scalable Online Recurrent Learning Using Columnar Neural Networks View paper
- [50] Learning to Control Rapidly Changing Synaptic Connections: An Alternative Type of Memory in Sequence Processing Artificial Neural Networks View paper
- [51] PredRNN: A Recurrent Neural Network for Spatiotemporal Predictive Learning View paper
- [52] Resource-Efficient Acoustic Full-Waveform Inversion via Dual-Branch Physics-Informed RNN with Scale Decomposition View paper
- [53] Predicting Wave Dynamics using Deep Learning with Multistep Integration Inspired Attention and Physics-Based Loss Decomposition View paper
- [54] Two-shot learning of continuous interpolation using a conceptor-aided recurrent autoencoder View paper
- [55] THE COMPUTATIONAL ROLE OF COMPLEX REPRESENTATIONS IN RNNS View paper
- [56] Predicting change points in multivariate time series data View paper
- [57] A temporal convolutional recurrent autoencoder based framework for compressing time series data View paper
- [58] Comparative study of state-based neural networks for virtual analog audio effects modeling View paper
- [59] HT-STNet: a hierarchical Tucker decomposition and spatio-temporal LSTM network for accurate and efficient shared mobility demand forecasting on sparse data View paper
- [60] Temporal superimposed crossover module for effective continuous sign language View paper
- [61] Recurrent neural networks for edge intelligence: a survey View paper
- [62] Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task View paper
- [63] An extended echo state network using Volterra filtering and principal component analysis View paper
- [64] Bbs-rnn: Block-based Structure Compression With Admm For Rnn On Temporal Sequence Applications. View paper
- [65] Fostering Event-Predictive Encodings in Recurrent Neural Networks View paper