Novelty Assessment Report

Paper: CIMemories: A Compositional Benchmark For Contextual Integrity In LLMs
PDF URL: https://openreview.net/pdf?id=YnNIp38v1M
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Large Language Models (LLMs) increasingly use persistent memory from past interactions to enhance personalization and task performance. However, this memory creates critical risks when sensitive information is revealed in inappropriate contexts. We present CIMemories, a benchmark for evaluating whether LLMs appropriately control information flow from memory based on task context. CIMemories uses synthetic user profiles with 100+ attributes per user, paired with various task contexts where each attribute may be essential for some tasks but inappropriate for others. For example, mental health details are necessary for booking therapy but inappropriate when requesting time off from work. This design enables two forms of compositionality: (1) flexible memory composition by varying which attributes are necessary versus inappropriate across different settings, and (2) multi-task composition per user, measuring cumulative information disclosure across sessions. Our evaluation reveals frontier models exhibit between 14%-69% attribute-level violations (leaking inappropriate information), and that higher task completeness (sharing necessary information) is accompanied by increased violations, highlighting critical gaps in integrity-aware memory systems.

Core Task Landscape

This paper addresses: **evaluating contextual information flow control in memory-augmented language models**
A total of **50 papers** were analyzed and organized into a taxonomy with **14 categories**.

Taxonomy Overview

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

- **Memory Architecture and Mechanisms**
- **Retrieval-Augmented Generation and Knowledge Integration**
- **Knowledge Conflict Resolution and Information Flow Control**
- **Contextual Integrity and Privacy-Aware Memory Systems**
- **Application-Specific Memory Systems**
- **Specialized Memory Mechanisms and Theoretical Frameworks**

Complete Taxonomy Tree

- evaluating contextual information flow control in memory-augmented language models Survey Taxonomy
- Memory Architecture and Mechanisms
  - Hierarchical and Multi-Tiered Memory Systems (9 papers)
  - [1] Cognitive memory in large language models (Shan LianLei, 2025) View paper
  - [4] Lightmem: Lightweight and efficient memory-augmented generation (Jizhan Fang, 2025) View paper
  - [13] Empowering working memory for large language model agents (Guo Jing, 2023) View paper
  - [16] Memory in Large Language Models: Mechanisms, Evaluation and Evolution (Zhang Dianxing, 2025) View paper
  - [19] LM2: Large Memory Models (Kang, 2025) View paper
  - [23] LM2: Large Memory Models for Long Context Reasoning (J Kang, 2025) View paper
  - [33] Semantic Hierarchical Reinforcement in Large Language Models for Contextual Memory Persistence (Nathan Poctik, 2025) View paper
  - [43] MMAG: Mixed Memory-Augmented Generation for Large Language Models Applications (Zeppieri, 2025) View paper
  - [48] Contextual Modulation Through Multi-Tiered Memory Networks in Large Language Models (Ruslan Stefanov, 2024) View paper
  - Recurrent and Sequential Memory Architectures (3 papers)
  - [28] Memory architectures in recurrent neural network language models (Dani Yogatama, 2018) View paper
  - [32] Can Memory-Augmented Language Models Generalize on Reasoning-in-a-Haystack Tasks? (Das, 2025) View paper
  - [37] Recurrent Memory-Augmented Transformers with Chunked Attention for Long-Context Language Modeling (Kashyap, 2025) View paper
  - Associative and Relational Memory Systems (3 papers)
  - [30] CAMELoT: Towards Large Language Models with Training-Free Consolidated Associative Memory (He, 2024) View paper
  - [35] Relational Memory-Augmented Language Models (Liu Qi, 2022) View paper
  - [38] TACL: Relational Memory Augmented Language Models (Blunsom, 2022) View paper
  - Context-Aware and Dynamic Memory Fusion (4 papers)
  - [2] Investigating contextual layer fusion in recent open source large language models for context retention and comprehension (Kristina Firstova, 2024) View paper
  - [12] Context-aware dynamic memory fusion in large language models for advanced task-specific performance (Vaughan Ikaris, 2024) View paper
  - [24] Exploring Synaptic Resonance in Large Language Models: A Novel Approach to Contextual Memory Integration (George Applegarth, 2025) View paper

- [44] Symbol-Rooted Cascade Propagation in Contextual Memory Routing for Large Language Models (Watson, 2025) View paper
- Memory-Augmented Attention and Embedding Mechanisms (2 papers)
- [11] Semantic coherence dynamics in large language models through layered syntax-aware memory retention mechanism (Carl Anderson, 2024) View paper
- [42] Enhancing Large Language Models through Dynamic Contextual Memory Embedding: A Technical Evaluation (Igor Dakat, 2024) View paper

- Retrieval-Augmented Generation and Knowledge Integration
  - Parametric-Nonparametric Knowledge Interaction (3 papers)
  - [3] Deciphering the Interplay of Parametric and Non-parametric Memory in Retrieval-augmented Language Models (Farahani, 2024) View paper
  - [5] Parameters vs. Context: Fine-Grained Control of Knowledge Reliance in Language Models (Liu Sheng-hua, 2025) View paper
  - [9] Adaptive semiparametric language models (Dani Yogatama, 2021) View paper
  - Long-Context and Memory-Enhanced Retrieval (4 papers)
  - [6] Memorag: Boosting long context processing with global memory-enhanced retrieval augmentation (Hongjin Qian, 2025) View paper
  - [8] Memory-Augmented Architecture for Long-Term Context Handling in Large Language Models (Usama, 2025) View paper
  - [26] Memlong: Memory-augmented retrieval for long text modeling (Liu Weijie, 2024) View paper
  - [46] Augmenting Language Models with Long-Term Memory (Wang, 2023) View paper
  - Retrieval-Augmented Reasoning and Planning (3 papers)
  - [15] MemoTime: Memory-Augmented Temporal Knowledge Graph Enhanced Large Language Model Reasoning (Tan Xing-yu, 2025) View paper
  - [17] RARE: Retrieval-Augmented Reasoning Enhancement for Large Language Models (Hieu Tran, 2024) View paper
  - [22] RAP: Retrieval-Augmented Planning with Contextual Memory for Multimodal LLM Agents (Kagaya Tomoyuki, 2024) View paper
  - Adaptive Retrieval and Context Compression (5 papers)
  - [14] Pretraining Context Compressor for Large Language Models with Embedding-Based Memory (Yuhong Dai, 2025) View paper
  - [18] Enhancing Cache-Augmented Generation (CAG) with Adaptive Contextual Compression for Scalable Knowledge Integration (Agrawal, 2025) View paper
  - [25] Online adaptation of language models with a memory of amortized contexts (Jaehyung Kim, 2024) View paper
  - [34] MeVe: A Modular System for Memory Verification and Effective Context Control in Language Models (Andreas Ottem, 2025) View paper
  - [50] NMRet: A Memory-Augmented Retrieval Framework for Large Language Models (S Bhat, n.d.) View paper
  - Multi-Dimensional and Multi-Modal Retrieval (3 papers)
  - [27] MIRA-CAP: Memory-Integrated Retrieval-Augmented Captioning for State-of-the-Art Image and Video Captioning (Sabina Umirzakova, 2024) View paper
  - [39] IMDMR: An Intelligent Multi-Dimensional Memory Retrieval System for Enhanced Conversational AI (Tejas Pawar, 2025) View paper
  - [41] Memory-Augmented Language Models (MALM): A Unified Architecture for Knowledge Integration and Safety Compliance (Li, 2025) View paper
- Knowledge Conflict Resolution and Information Flow Control (4 papers)
  - [10] Cutting Off the Head Ends the Conflict: A Mechanism for Interpreting and Mitigating Knowledge Conflicts in Language Models (Zhuoran Jin, 2024) View paper
  - [29] Studying Large Language Model Behaviors Under Context-Memory Conflicts With Real Documents (Kortukov, 2024) View paper
  - [31] Do Retrieval-Augmented Language Models Adapt to Varying User Needs? (Wu Peilin, 2025) View paper
  - [47] Large Language Models with Controllable Working Memory (Daliang Li, 2023) View paper
- Contextual Integrity and Privacy-Aware Memory Systems ★ (3 papers)
  - [0] CIMemories: A Compositional Benchmark For Contextual Integrity In LLMs (Anon et al., 2026) View paper
  - [20] Memos: A memory os for ai system (Li Zhiyu, 2025) View paper
  - [40] CIMemories: A Compositional Benchmark for Contextual Integrity of Persistent Memory in LLMs (Niloofar Mireshghallah, 2025) View paper
- Application-Specific Memory Systems (3 papers)
  - [36] Natural Language Processing-Driven Document Summarization Using Attention-Guided Memory-Augmented Transformer (Suresh Kurapati, 2025) View paper
  - [45] Better conditioning on context for Natural Language Processing (Qi, 2022) View paper
  - [49] Memory Augmented Lookup Dictionary based Language Modeling for Automatic Speech Recognition (Yukun Feng, 2022) View paper
- Specialized Memory Mechanisms and Theoretical Frameworks (2 papers)
  - [7] Fractal lattice conditioning for structured divergence in large language model gradient spaces (E Littlewood, 2025) View paper
  - [21] Understanding intra-node communication in HPC systems and Datacenters (J Tarraga-Moreno, 2025) View paper

Narrative

Core task: evaluating contextual information flow control in memory-augmented language models. The field has evolved into several interconnected branches that address how language models store, retrieve, and manage information. Memory Architecture and Mechanisms explores foundational designs—ranging from working memory systems like Working Memory Agents[13] to multi-tiered structures such as Multi Tiered Memory[48]—that determine how models organize and access stored knowledge. Retrieval-Augmented Generation and Knowledge Integration focuses on methods that blend parametric model knowledge with external retrieval, exemplified by works like Memorag[6] and Cache Augmented Generation[18], which balance efficiency and accuracy when incorporating retrieved context. Knowledge Conflict Resolution and Information Flow Control tackles the challenge of reconciling contradictory information from different sources, as seen in Context Memory Conflicts[29] and Parametric Nonparametric Memory[3]. Meanwhile, Contextual Integrity and Privacy-Aware Memory Systems emphasizes controlling what information flows where, ensuring that sensitive data respects privacy boundaries and contextual norms. Application-Specific Memory Systems tailors memory mechanisms to domains like dialogue or code generation, and Specialized Memory Mechanisms and Theoretical Frameworks investigates novel architectures such as Recurrent Memory Transformers[37] and theoretical underpinnings of memory dynamics.

A particularly active line of work examines the trade-offs between parametric storage and dynamic retrieval: some studies like Parameters vs Context[5] and Lightmem[4] investigate when to rely on model weights versus external memory, while others such as Adaptive Semiparametric[9] propose hybrid strategies. Another emerging theme is the need for principled information flow control,

especially when memory systems must respect privacy or contextual boundaries. CIMemories[0] sits squarely within the Contextual Integrity and Privacy-Aware Memory Systems branch, addressing how to evaluate whether memory-augmented models properly enforce contextual norms when retrieving and using stored information. Its emphasis on formal evaluation of information flow distinguishes it from neighbors like Memory OS[20], which focuses more on system-level memory management, and CIMemories Persistent[40], which extends similar ideas to persistent storage scenarios. Together, these works highlight an open question: how can we rigorously verify that memory mechanisms respect the intended boundaries of information sharing in complex, multi-context environments?

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. Memos: A memory os for ai system

Authors: Li Zhiyu, Song Shi-chao, Zhiyu Li, Xi, Chenyang, et al. (81 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Large Language Models (LLMs) have become an essential infrastructure for Artificial General Intelligence (AGI), yet their lack of well-defined memory management systems hinders the development of long-context reasoning, continual personalization, and knowledge consistency.Existing models mainly rely on static parameters and short-lived contextual states, limiting their ability to track user preferences or update knowledge over extended periods.While Retrieval-Augmented Generation (RAG) introduce...

Relationship Analysis

Both papers belong to the category of Contextual Integrity and Privacy-Aware Memory Systems, focusing on how memory-augmented LLMs handle information disclosure. While the original paper (CIMemories) provides a compositional benchmark for evaluating contextual integrity violations when LLMs retrieve and share personal attributes across different social contexts, the candidate paper (MemOS) proposes a memory operating system architecture that manages memory lifecycle, scheduling, and evolution across plaintext, activation, and parameter-level representations. The key difference is that CIMemories is an evaluation framework measuring privacy violations, whereas MemOS is a system design for memory management that could potentially be evaluated using such benchmarks.

---

2. CIMemories: A Compositional Benchmark for Contextual Integrity of Persistent Memory in LLMs

Authors: Niloofar Mireshghallah, Neal Mangaokar, Narine Kokhlikyan, Arman Zharmagambetov, Manzil Zaheer, et al. (7 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Large Language Models (LLMs) increasingly use persistent memory from past interactions to enhance personalization and task performance. However, this memory introduces critical risks when sensitive information is revealed in inappropriate contexts. We present CIMemories, a benchmark for evaluating whether LLMs appropriately control information flow from memory based on task context. CIMemories uses synthetic user profiles with over 100 attributes per user, paired with diverse task contexts in wh...

⚠ Similarity Notice

This paper appears to be the same work as the original paper or a very close variant. Both papers introduce CIMemories as a benchmark for evaluating contextual integrity in memory-augmented LLMs, use identical synthetic user profiles with 100+ attributes, employ the same evaluation framework measuring violations and completeness, and report nearly identical experimental findings (e.g., GPT-4o at ~14% violations, Qwen-3 32B at ~69% violations). The core methodology, dataset construction approach, and key results are essentially identical across both papers.

Contributions Analysis

Overall novelty summary. The paper introduces CIMemories, a benchmark for evaluating whether memory-augmented LLMs appropriately control information disclosure across different task contexts. It resides in the 'Contextual Integrity and Privacy-Aware Memory Systems' leaf, which contains only three papers total, making this a relatively sparse research direction within the broader taxonomy of 50 papers. The sibling papers include work on memory operating systems and persistent memory scenarios, but none directly address the compositional evaluation of contextual information flow that CIMemories targets.

The taxonomy reveals that most memory research concentrates on architectural designs (nine papers in hierarchical systems alone) and retrieval-augmented generation (multiple subcategories with 15+ papers). The neighboring 'Knowledge Conflict Resolution and Information Flow Control' branch addresses contradictions between parametric and contextual knowledge but does not emphasize privacy-aware disclosure control. CIMemories bridges a gap between these architectural concerns and the emerging need for principled information flow evaluation, connecting privacy considerations to the broader memory-augmented LLM ecosystem.

Among 22 candidates examined across three contributions, none were found to clearly refute the paper's claims. The core benchmark contribution examined 10 candidates with zero refutable matches, the compositional design examined 10 candidates with zero refutable matches, and the privacy persona labeling method examined 2 candidates with zero refutable matches. This suggests that within the limited search scope, the specific combination of contextual integrity evaluation, compositional memory design, and scalable labeling appears relatively unexplored in prior work.

Based on the top-22 semantic matches examined, the work appears to occupy a novel position at the intersection of memory systems and privacy-aware information flow. The sparse population of its taxonomy leaf and absence of refuting candidates within the search scope suggest substantive novelty, though a broader literature search might reveal additional related efforts in privacy-preserving NLP or access control systems not captured by this memory-focused taxonomy.

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

Contribution 1: CIMemories benchmark for evaluating contextual integrity in memory-augmented LLMs

Description: The authors introduce CIMemories, a benchmark that uses synthetic user profiles with over 100 attributes per user paired with various task contexts to evaluate whether LLMs respect contextual integrity when using persistent memory. The benchmark enables compositional evaluation through flexible memory composition and multi-task composition per user.

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. Private Memorization Editing: Turning Memorization into a Defense to Strengthen Data Privacy in Large Language Models

URL: View paper

Brief Assessment

Private Memorization Editing[62] focuses on preventing memorized PII leakage through model editing techniques, not on evaluating contextual integrity in memory-augmented systems. The candidate addresses privacy through parameter editing of memorized training data, while the original contribution evaluates whether LLMs appropriately control information flow based on task context.

---

2. Preventing generation of verbatim memorization in language models gives a false sense of privacy

URL: View paper

Brief Assessment

Verbatim Memorization[61] focuses on preventing verbatim and approximate memorization of training data in language models, not on evaluating contextual integrity in memory-augmented systems with user profiles and task contexts.

---

3. Preserving privacy through dememorization: An unlearning technique for mitigating memorization risks in language models

URL: View paper

Brief Assessment

Dememorization[69] addresses memorization of training data through unlearning techniques, not contextual integrity evaluation in memory-augmented systems. The candidate focuses on preventing verbatim reproduction of training data, while the original evaluates appropriate information flow across different social contexts.

---

4. Dynamic semantic memory retention in large language models: An exploration of spontaneous retrieval mechanisms

URL: View paper

Brief Assessment

Semantic Memory Retention[67] focuses on memory retention mechanisms and gating functions for information flow across sequences, not on evaluating contextual integrity or privacy violations in memory-augmented systems.

---

5. Dynamic neural alignment mechanisms in large language models to contextual integrity preservation

URL: View paper

Brief Assessment

Neural Alignment Mechanisms[68] focuses on neural component mechanisms for contextual integrity preservation in LLMs, not on benchmark development for evaluating memory-augmented systems with compositional user profiles and task contexts.

---

6. Effectiveness of Privacy-preserving Algorithms in LLMs: A Benchmark and Empirical Analysis

URL: View paper

Brief Assessment

Privacy Preserving Benchmark[64] focuses on evaluating privacy-preserving algorithms (like differential privacy) for protecting training data and user queries in LLMs, not on contextual integrity in memory-augmented systems where the same information may be appropriate in some contexts but inappropriate in others.

---

7. An LLM-enabled human demonstration-assisted hybrid robot skill synthesis approach for human-robot collaborative assembly

URL: View paper

Brief Assessment

Human Demonstration Assembly[65] focuses on robot skill synthesis using LLM-enabled learning from demonstration and deep reinforcement learning for human-robot collaborative assembly tasks, not on evaluating contextual integrity or privacy in memory-augmented language models.

---

8. Transformer-based generative memory embedding for adaptive contextual recall

URL: View paper

Brief Assessment

Generative Memory Embedding[66] focuses on memory embedding techniques for adaptive contextual recall, not on benchmarking contextual integrity violations in memory-augmented systems. The candidate's limited context does not demonstrate prior work on compositional evaluation of contextual integrity with synthetic user profiles.

---

9. Advancing Conversational Psychotherapy: Integrating Privacy, Dual-Memory, and Domain Expertise with Large Language Models

URL: View paper

Brief Assessment

Conversational Psychotherapy[63] focuses on psychotherapy chatbots with dual-memory for personalized therapy responses and privacy preservation in clinical contexts, not on evaluating contextual integrity across diverse social contexts with compositional benchmarking.

---

10. Memory in Large Language Models: Mechanisms, Evaluation and Evolution

URL: View paper

Brief Assessment

Memory Mechanisms Evaluation[16] focuses on a general taxonomy and evaluation framework for LLM memory mechanisms (parametric, contextual, external, procedural/episodic), not on contextual integrity or privacy-preserving information flow in specific social contexts.

Contribution 2: Compositional design with flexible memory and multi-task composition

Description: The benchmark features a novel compositional design that allows dynamic variation of which attributes are necessary versus inappropriate across different settings, and measures cumulative information disclosure across multiple tasks per user to study how violations accumulate over time.

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. Improving Deep Learning Framework Testing with Model-Level Metamorphic Testing

URL: View paper

Brief Assessment

Model Level Metamorphic[55] focuses on testing deep learning frameworks through metamorphic relations on model structures, not on evaluating information disclosure across multiple tasks with memory configurations in LLM assistants.

---

2. The sum leaks more than its parts: Compositional privacy risks and mitigations in multi-agent collaboration

URL: View paper

Brief Assessment

Compositional Privacy Risks[53] focuses on compositional privacy leakage across multiple agents in distributed systems, where sensitive information emerges from combining outputs from different agents. The original paper's compositional design varies memory attributes and evaluates cumulative disclosure across tasks for a single user, which is a different evaluation paradigm than multi-agent compositional inference.

---

3. Retaining privileged information for multi-task learning

URL: View paper

Brief Assessment

Retaining Privileged Information[57] focuses on multi-task learning through privileged information transfer in a teacher-student framework, not on evaluating information disclosure or contextual integrity across tasks with varying memory compositions.

---

4. Context Parametrization with Compositional Adapters

URL: View paper

Brief Assessment

Compositional Adapters[59] focuses on compositional adapter generation for LLM context parametrization (demonstrations, instructions, passages), not on compositional evaluation of information disclosure across memory configurations and multiple user tasks as in the original paper's privacy benchmark.

---

5. A cut principle for information flow

URL: View paper

Brief Assessment

Cut Principle[60] focuses on information flow constraints in distributed systems through graph structure and cut sets, not on compositional evaluation of memory configurations or multi-task information disclosure in LLM contexts.

---

6. CIMemories: A Compositional Benchmark for Contextual Integrity of Persistent Memory in LLMs

URL: View paper

Brief Assessment

CIMemories Persistent[40] appears to be the same work as the original paper, using identical methodology and evaluation framework. The candidate text describes the same benchmark with the same compositional design features.

---

7. User profiling and satisfaction inference in public information access services

URL: View paper

Brief Assessment

User Profiling Satisfaction[56] focuses on user profiling and satisfaction inference in public information access services, not on compositional evaluation of information disclosure across multiple tasks with memory configurations in LLM systems.

---

8. Strategic information provision in multidimensional environments

URL: View paper

Brief Assessment

Strategic Information Provision[58] is a doctoral thesis in economic sciences focused on strategic information provision in multidimensional environments. The provided context contains only the title page and does not reveal technical content about compositional evaluation frameworks or memory systems that would refute the original paper's novelty claims about benchmark design for contextual integrity in LLMs.

---

9. Memory in Large Language Models: Mechanisms, Evaluation and Evolution

URL: View paper

Brief Assessment

Memory Mechanisms Evaluation[16] addresses temporal governance and cross-session consistency in memory systems, but does not propose compositional evaluation of information disclosure across varying attribute necessity or cumulative privacy violations across multiple tasks.

---

10. The Routledge Handbook of Behavioural Accounting Research

URL: View paper

Brief Assessment

Behavioural Accounting Research[54] appears to be a handbook on behavioral accounting topics (audit evidence, performance measures) and does not address compositional evaluation of LLM memory systems or information disclosure across multiple tasks.

Contribution 3: Scalable contextual integrity labeling using privacy personas

Description: The authors develop a scalable method for generating contextual integrity ground truth labels by using multiple privacy personas from established surveys, sampling labels multiple times per persona, and assigning final labels only where all personas agree, thereby respecting the inherent subjectivity in privacy norms while enabling large-scale evaluation.

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. Search with discretion: value sensitive design of training data for information retrieval

URL: View paper

Brief Assessment

Search with Discretion[51] mentions using 'values personas' for annotation, but the provided context is too fragmentary to determine if this approach predates or refutes the original paper's method of using multiple privacy personas with agreement-based labeling for contextual integrity.

---

2. Privacy Awareness for Information-Sharing Assistants: A Case-study on Form-filling with Contextual Integrity

URL: View paper

Brief Assessment

Privacy Awareness Assistants[52] uses privacy personas for form-filling tasks with human annotations, while the original paper applies personas to generate CI labels for memory-augmented LLM benchmarks. The application domains and technical implementations differ substantially.

Appendix: Text Similarity Detection

Textual similarity detection checked 22 papers and found 1 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. CIMemories: A Compositional Benchmark for Contextual Integrity of Persistent Memory in LLMs

Detected in: Core Task (sibling), Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] CIMemories: A Compositional Benchmark For Contextual Integrity In LLMs View paper
- [1] Cognitive memory in large language models View paper
- [2] Investigating contextual layer fusion in recent open source large language models for context retention and comprehension View paper
- [3] Deciphering the Interplay of Parametric and Non-parametric Memory in Retrieval-augmented Language Models View paper
- [4] Lightmem: Lightweight and efficient memory-augmented generation View paper
- [5] Parameters vs. Context: Fine-Grained Control of Knowledge Reliance in Language Models View paper
- [6] Memorag: Boosting long context processing with global memory-enhanced retrieval augmentation View paper
- [7] Fractal lattice conditioning for structured divergence in large language model gradient spaces View paper
- [8] Memory-Augmented Architecture for Long-Term Context Handling in Large Language Models View paper
- [9] Adaptive semiparametric language models View paper
- [10] Cutting Off the Head Ends the Conflict: A Mechanism for Interpreting and Mitigating Knowledge Conflicts in Language Models View paper
- [11] Semantic coherence dynamics in large language models through layered syntax-aware memory retention mechanism View paper
- [12] Context-aware dynamic memory fusion in large language models for advanced task-specific performance View paper
- [13] Empowering working memory for large language model agents View paper
- [14] Pretraining Context Compressor for Large Language Models with Embedding-Based Memory View paper
- [15] MemoTime: Memory-Augmented Temporal Knowledge Graph Enhanced Large Language Model Reasoning View paper
- [16] Memory in Large Language Models: Mechanisms, Evaluation and Evolution View paper
- [17] RARE: Retrieval-Augmented Reasoning Enhancement for Large Language Models View paper
- [18] Enhancing Cache-Augmented Generation (CAG) with Adaptive Contextual Compression for Scalable Knowledge Integration View paper
- [19] LM2: Large Memory Models View paper
- [20] Memos: A memory os for ai system View paper
- [21] Understanding intra-node communication in HPC systems and Datacenters View paper
- [22] RAP: Retrieval-Augmented Planning with Contextual Memory for Multimodal LLM Agents View paper
- [23] LM2: Large Memory Models for Long Context Reasoning View paper
- [24] Exploring Synaptic Resonance in Large Language Models: A Novel Approach to Contextual Memory Integration View paper
- [25] Online adaptation of language models with a memory of amortized contexts View paper
- [26] Memlong: Memory-augmented retrieval for long text modeling View paper
- [27] MIRA-CAP: Memory-Integrated Retrieval-Augmented Captioning for State-of-the-Art Image and Video Captioning View paper
- [28] Memory architectures in recurrent neural network language models View paper
- [29] Studying Large Language Model Behaviors Under Context-Memory Conflicts With Real Documents View paper
- [30] CAMELoT: Towards Large Language Models with Training-Free Consolidated Associative Memory View paper
- [31] Do Retrieval-Augmented Language Models Adapt to Varying User Needs? View paper
- [32] Can Memory-Augmented Language Models Generalize on Reasoning-in-a-Haystack Tasks? View paper
- [33] Semantic Hierarchical Reinforcement in Large Language Models for Contextual Memory Persistence View paper
- [34] MeVe: A Modular System for Memory Verification and Effective Context Control in Language Models View paper
- [35] Relational Memory-Augmented Language Models View paper
- [36] Natural Language Processing-Driven Document Summarization Using Attention-Guided Memory-Augmented Transformer View paper
- [37] Recurrent Memory-Augmented Transformers with Chunked Attention for Long-Context Language Modeling View paper
- [38] TACL: Relational Memory Augmented Language Models View paper
- [39] IMDMR: An Intelligent Multi-Dimensional Memory Retrieval System for Enhanced Conversational AI View paper
- [40] CIMemories: A Compositional Benchmark for Contextual Integrity of Persistent Memory in LLMs View paper
- [41] Memory-Augmented Language Models (MALM): A Unified Architecture for Knowledge Integration and Safety Compliance View paper
- [42] Enhancing Large Language Models through Dynamic Contextual Memory Embedding: A Technical Evaluation View paper
- [43] MMAG: Mixed Memory-Augmented Generation for Large Language Models Applications View paper
- [44] Symbol-Rooted Cascade Propagation in Contextual Memory Routing for Large Language Models View paper
- [45] Better conditioning on context for Natural Language Processing View paper
- [46] Augmenting Language Models with Long-Term Memory View paper
- [47] Large Language Models with Controllable Working Memory View paper
- [48] Contextual Modulation Through Multi-Tiered Memory Networks in Large Language Models View paper
- [49] Memory Augmented Lookup Dictionary based Language Modeling for Automatic Speech Recognition View paper
- [50] NMRet: A Memory-Augmented Retrieval Framework for Large Language Models View paper
- [51] Search with discretion: value sensitive design of training data for information retrieval View paper

- [52] Privacy Awareness for Information-Sharing Assistants: A Case-study on Form-filling with Contextual Integrity View paper
- [53] The sum leaks more than its parts: Compositional privacy risks and mitigations in multi-agent collaboration View paper
- [54] The Routledge Handbook of Behavioural Accounting Research View paper
- [55] Improving Deep Learning Framework Testing with Model-Level Metamorphic Testing View paper
- [56] User profiling and satisfaction inference in public information access services View paper
- [57] Retaining privileged information for multi-task learning View paper
- [58] Strategic information provision in multidimensional environments View paper
- [59] Context Parametrization with Compositional Adapters View paper
- [60] A cut principle for information flow View paper
- [61] Preventing generation of verbatim memorization in language models gives a false sense of privacy View paper
- [62] Private Memorization Editing: Turning Memorization into a Defense to Strengthen Data Privacy in Large Language Models View paper
- [63] Advancing Conversational Psychotherapy: Integrating Privacy, Dual-Memory, and Domain Expertise with Large Language Models View paper
- [64] Effectiveness of Privacy-preserving Algorithms in LLMs: A Benchmark and Empirical Analysis View paper
- [65] An LLM-enabled human demonstration-assisted hybrid robot skill synthesis approach for human-robot collaborative assembly View paper
- [66] Transformer-based generative memory embedding for adaptive contextual recall View paper
- [67] Dynamic semantic memory retention in large language models: An exploration of spontaneous retrieval mechanisms View paper
- [68] Dynamic neural alignment mechanisms in large language models to contextual integrity preservation View paper
- [69] Preserving privacy through dememorization: An unlearning technique for mitigating memorization risks in language models View paper