

# Novelty Assessment Report

**Paper:** Context Parametrization with Compositional Adapters

**PDF URL:** <https://openreview.net/pdf?id=8oRi7CEj6A>

**Venue:** ICLR 2026 Conference Submission

**Year:** 2026

**Report Generated:** 2026-01-08

## Abstract

Large language models (LLMs) often seamlessly adapt to new tasks through in-context learning (ICL) or supervised fine-tuning (SFT). However, both of these approaches face key limitations: ICL is inefficient when handling many demonstrations, and SFT incurs training overhead while sacrificing flexibility.

Mapping instructions or demonstrations from context directly into adapter parameters offers an appealing alternative. While prior work explored generating adapters based on a single input context, it has overlooked the need to integrate multiple chunks of information. To address this gap, we introduce CompAs, a meta-learning framework that translates context into adapter parameters with a compositional structure.

Adapters generated this way can be merged algebraically, enabling instructions, demonstrations, or retrieved passages to be seamlessly combined without reprocessing long prompts.

Critically, this approach yields three benefits: lower inference cost, robustness to long-context instability, and establishes a principled solution when input exceeds the model's context window.

Furthermore, CompAs encodes information into adapter parameters in a reversible manner, enabling recovery of input context through a decoder, facilitating safety and security. Empirical results on diverse multiple-choice and extractive question answering tasks show that \method outperforms ICL and prior generator-based methods, especially when scaling to more inputs. Our work establishes composable adapter generation as a practical and efficient alternative for scaling LLM deployment.

### Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: [mingzhang23@m.fudan.edu.cn](mailto:mingzhang23@m.fudan.edu.cn)

## Core Task Landscape

This paper addresses: **Generating Compositional Adapter Parameters from Context for Language Models**

A total of **10 papers** were analyzed and organized into a taxonomy with **9 categories**.

### Taxonomy Overview

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

- **Context-to-Adapter Generation Methods**
- **Task-Aware and Context-Oriented Adapter Design**
- **Multi-Task Adapter Composition and Fusion**
- **Specialized Adapter Applications and Architectures**

### Complete Taxonomy Tree

- Generating Compositional Adapter Parameters from Context for Language Models Survey Taxonomy
- Context-to-Adapter Generation Methods
  - Single-Pass Generative Adapter Synthesis (1 papers)
  - [5] Generative adapter: Contextualizing language models in parameters with a single forward pass (Chen Tong, 2024) [View paper](#)
  - Compositional Multi-Context Adapter Generation ★ (1 papers)
  - [0] Context Parametrization with Compositional Adapters (Anon et al., 2026) [View paper](#)
- Task-Aware and Context-Oriented Adapter Design
  - Weight Decomposition with Task Context (2 papers)
  - [2] Corda: Context-oriented decomposition adaptation of large language models for task-aware parameter-efficient fine-tuning (Yi-bo, 2024) [View paper](#)
  - [3] CorDA: Context-Oriented Decomposition Adaptation of Large Language Models (Yang Yi-bo, 2024) [View paper](#)
  - Graph-Structured Multi-Task Adapter Composition (1 papers)
  - [1] Structural Priors and Modular Adapters in the Composable Fine-Tuning Algorithm of Large-Scale Models (Wang Yuxiao, 2025) [View paper](#)
- Multi-Task Adapter Composition and Fusion
  - Two-Stage Knowledge Extraction and Composition (1 papers)
  - [4] Adapterfusion: Non-destructive task composition for transfer learning (Jonas Pfeiffer, 2021) [View paper](#)
  - Conditional and Compositional Prompt-Adapter Systems (1 papers)
  - [8] On conditional and compositional language model differentiable prompting (Jonathan Pilault, 2023) [View paper](#)
- Specialized Adapter Applications and Architectures
  - Compositional Zero-Shot Learning with Intra-Layer Adapters (1 papers)
  - [7] Caila: concept-aware intra-layer adapters for compositional zero-shot learning (Zhaoheng Zheng, 2024) [View paper](#)
  - Domain-Specific Adapter Integration (2 papers)
  - [6] Latent symbol lattices in probabilistic semiosis: An unconventional architectural mechanism for contextual modulation in large language models (R Marchbanks, 2025) [View paper](#)
  - [10] A Context Awareness Framework for Real-Time Analysis and Adaptive Navigation in Dynamic Indoor Environments (R Awad, n.d.) [View paper](#)

- Modular Adapter Development Frameworks (1 papers)
- [9] AdapterHub for Modular Natural Language Processing: The Complete Guide for Developers and Engineers (Smith, 2025) [View paper](#)

## Narrative

Core task: generating compositional adapter parameters from context for language models. The field centers on making language models more flexible by dynamically producing or combining adapter modules based on input context, rather than relying solely on fixed, pre-trained adapters. The taxonomy reveals four main branches. Context-to-Adapter Generation Methods explore techniques that directly map contextual signals into adapter weights, enabling on-the-fly parameterization. Task-Aware and Context-Oriented Adapter Design focuses on crafting adapters that are sensitive to task-specific or contextual cues, often leveraging structural priors or domain knowledge. Multi-Task Adapter Composition and Fusion investigates how to merge or orchestrate multiple adapters—such as AdapterFusion[4]—to handle diverse tasks simultaneously. Finally, Specialized Adapter Applications and Architectures address domain-specific use cases and novel architectural variants, demonstrating the breadth of adapter-based approaches across different problem settings.

Recent work has intensified around compositional and generative strategies. Several studies, including Corda[2], CorDA[3], and Generative Adapter[5], emphasize learning to produce adapter parameters conditioned on input or task context, trading off between expressiveness and computational overhead. Compositional Adapters[0] sits squarely within this line of inquiry, focusing on multi-context scenarios where adapters must be generated and composed dynamically. It shares conceptual ground with CorDA[3], which also targets context-driven adapter generation, but Compositional Adapters[0] places stronger emphasis on handling multiple contextual signals simultaneously. Meanwhile, approaches like Structural Priors Adapters[1] and Caila[7] explore how to inject inductive biases or task-aware structures into adapter design, offering a complementary perspective on context sensitivity. Together, these directions highlight ongoing questions about scalability, interpretability, and the optimal granularity for context-driven parameterization.

## Related Works in Same Category

---

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

### Taxonomy-Level Summary

Both subtopics address the challenge of generating adapter parameters from context for language models, representing different architectural approaches to context-driven adaptation. The original leaf focuses on compositional methods that handle multiple context chunks through algebraic merging operations, while its sibling handles single-context scenarios with direct generation. These represent complementary strategies along the dimension of context complexity and composition mechanisms.

**Similarities:** - Both generate adapter parameters dynamically from input context rather than learning fixed adapters - Both aim to enable rapid adaptation of language models without extensive fine-tuning - Both leverage context information to produce task-specific or domain-specific model modifications - Both fall under the broader paradigm of context-conditioned parameter generation

**Differences:** - Original leaf handles multiple context chunks requiring algebraic composition/merging, while sibling processes single context in one pass - Original leaf emphasizes compositional operations (algebraic merging) as a core mechanism, while sibling focuses on direct generation efficiency - Original leaf enables seamless integration of information from disparate sources, while sibling optimizes for simplicity and speed with unified context - Sibling explicitly mentions self-supervised or meta-learning training paradigms, while original leaf emphasizes the compositional architecture itself

**Suggested Search Directions:** - Investigate hybrid approaches that combine single-pass efficiency with compositional capabilities for variable-length contexts - Explore the trade-offs between compositional complexity and generation quality/speed - Examine how different algebraic composition operators (addition, concatenation, attention-based merging) affect multi-context adapter performance

### Sibling Subtopics

- **Single-Pass Generative Adapter Synthesis** (leaves: 1, papers: 1)
- Scope: Approaches generating adapter parameters from context in a single forward pass using self-supervised or meta-learning.
- Exclude: Excludes compositional multi-context methods requiring algebraic merging; see Compositional Multi-Context Adapter Generation.

## Contributions Analysis

---

**Overall novelty summary.** The paper introduces CompAs, a meta-learning framework that generates adapter parameters from multiple context chunks and enables algebraic composition of these adapters. According to the taxonomy, this work resides in the 'Compositional Multi-Context Adapter Generation' leaf under 'Context-to-Adapter Generation Methods'. Notably, this leaf contains only the original paper itself—no sibling papers are listed. The parent category 'Context-to-Adapter Generation Methods' contains just one other leaf ('Single-Pass Generative Adapter Synthesis'), suggesting this is a relatively sparse research direction within the broader adapter landscape.

The taxonomy reveals that the broader field encompasses four main branches: context-to-adapter generation, task-aware adapter design, multi-task composition, and specialized applications. The paper's approach bridges context-to-adapter generation with multi-task composition concerns, as it addresses how to integrate multiple information sources without reprocessing. Neighboring work in 'Task-Aware and Context-Oriented Adapter Design' focuses on structural priors and task decomposition, while 'Multi-Task Adapter Composition and Fusion' explores combining pre-trained adapters. CompAs diverges by generating composable adapters on-the-fly rather than fusing pre-existing modules, positioning it at a distinct methodological intersection.

Among the three contributions analyzed, the literature search examined 30 candidate papers total. The core CompAs framework and theoretical composition conditions each examined 10 candidates with zero refutable prior work identified. The reversible encoding contribution examined 10 candidates and found 1 that appears to provide overlapping prior work. This suggests the compositional generation mechanism and theoretical foundations represent relatively unexplored territory within the limited search scope, while the reversibility aspect has at least some precedent. The analysis explicitly notes this is based on top-K semantic search plus citation expansion, not exhaustive coverage.

Given the limited search scope of 30 candidates and the sparse taxonomy leaf containing only this paper, the work appears to occupy a relatively novel position within context-driven adapter generation. However, the single-paper leaf status and modest search scale mean substantial related work may exist outside the examined candidates. The reversibility finding indicates at least one dimension has prior exploration, warranting careful positioning against that specific precedent.

---

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

### Contribution 1: COMPAS meta-learning framework for compositional adapter generation

**Description:** The authors propose COMPAS, a teacher-student framework that maps contextual information (instructions, demonstrations, or retrieved passages) into adapter parameters that can be algebraically merged. This enables seamless combination of multiple information sources without reprocessing long prompts, addressing efficiency and long-context instability issues.

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. Meta-LoRA: Meta-Learning LoRA Components for Domain-Aware ID Personalization

URL: [View paper](#)

### Brief Assessment

Meta LoRA[33] focuses on identity personalization in text-to-image models using meta-learned LoRA components for domain-specific priors, not on generating compositional adapters from contextual information (instructions, demonstrations, or retrieved passages) for language model adaptation as in COMPAS.

---

## 2. Modular meta-learning

URL: [View paper](#)

### Brief Assessment

Modular Meta Learning[40] focuses on learning reusable neural network modules that can be structurally recombined for new tasks through simulated annealing search, rather than generating adapter parameters from contextual information like instructions or demonstrations. The candidate addresses structural composition of learned modules across tasks, while COMPAS specifically maps context (instructions, demonstrations, passages) into adapter parameters with algebraic merging properties.

---

## 3. Adaptive compositional continual meta-learning

URL: [View paper](#)

### Brief Assessment

Compositional Continual Meta[36] focuses on continual meta-learning with compositional meta-knowledge for handling sequential heterogeneous tasks, not on generating adapter parameters from contextual information (instructions, demonstrations, or retrieved passages) for in-context learning as COMPAS does.

---

## 4. Neural Relational Inference with Fast Modular Meta-learning

URL: [View paper](#)

### Brief Assessment

Fast Modular Meta[39] focuses on meta-learning modular graph neural networks for relational inference in dynamical systems, not on generating compositional adapter parameters from contextual information for language models. The domains and technical approaches are fundamentally different.

---

## 5. MDANet: A multi-stage domain adaptation framework for generalizable low-light image enhancement

URL: [View paper](#)

### Brief Assessment

MDANet[38] focuses on domain adaptation for low-light image enhancement using multi-task meta-learning for visual tasks, not on generating compositional adapter parameters from textual contextual information (instructions, demonstrations, or passages) for language models.

---

## 6. Neural-Fly enables rapid learning for agile flight in strong winds

URL: [View paper](#)

### Brief Assessment

Neural Fly[32] focuses on learning aerodynamic representations for quadrotor flight control in wind conditions, not on generating compositional adapter parameters from contextual information for language models. The domains and technical approaches are fundamentally different.

---

## 7. Human-like systematic generalization through a meta-learning neural network

URL: [View paper](#)

### Brief Assessment

Systematic Generalization[31] focuses on meta-learning for compositional language understanding through instruction learning tasks, not on generating adapter parameters from contextual information for LLM adaptation. The candidate addresses systematic compositional generalization in cognitive modeling, while COMPAS addresses efficient context parametrization for LLMs.

---

## 8. Meta-learning hyperparameters for foundation model adaptation in remote-sensing imagery

URL: [View paper](#)

### Brief Assessment

Meta Learning Hyperparameters[37] focuses on meta-learning hyperparameters (ranks, dimensions) for foundation model adaptation in remote-sensing imagery, not on generating compositional adapters from contextual information that can be algebraically merged.

---

## 9. TAML-Adapter: Enhancing Adapter Tuning Through Task-Agnostic Meta-Learning for Low-Resource Automatic Speech Recognition

URL: [View paper](#)

### Brief Assessment

TAML-Adapter[35] focuses on meta-learning for adapter initialization in low-resource automatic speech recognition, not on compositional adapter generation from contextual information like instructions or demonstrations for language models.

---

## 10. AdaFML: Adaptive Federated Meta Learning With Multi-Objectives and Context-Awareness in Dynamic Heterogeneous Networks

URL: [View paper](#)

### Brief Assessment

AdaFML[34] focuses on federated learning with device-level context (data distribution, computation, communication conditions) for distributed model training, not on generating compositional adapter parameters from textual contextual information (instructions, demonstrations, passages) for single-model adaptation.

---

## Contribution 2: Theoretical conditions for parameter-space composition

**Description:** The authors formalize compositionality requirements through a monoid homomorphism framework and prove a compositionality bound (Theorem 1) that decomposes student-teacher error into generator additivity error and misfit on concatenated contexts, providing theoretical guarantees for when adapter addition approximates context concatenation.

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. Theory of overparametrization in quantum neural networks

URL: [View paper](#)

### Brief Assessment

Overparametrization Quantum Networks[21] focuses on quantum neural networks and their parameter optimization through dynamical Lie algebra theory, not on compositional adapter generation or parameter-space composition for context integration in language models.

---

## 2. Theoretical Investigation of Composite Neural Network

URL: [View paper](#)

### Brief Assessment

Composite Neural Network[27] focuses on combining pre-trained neural network models in a directed acyclic graph structure with frozen weights, proving performance bounds for ensemble-like compositions. The ORIGINAL paper addresses adapter generation from context with algebraic composition guarantees via monoid homomorphisms, which is a fundamentally different problem domain (context parametrization vs. model ensemble).

---

## 3. Barron Spaces and the Compositional Function Spaces for Neural Network Models

URL: [View paper](#)

### Brief Assessment

Barron Compositional Spaces[28] focuses on compositional function spaces for neural networks through flow-induced norms and approximation theory, not on adapter composition or parameter-space operations for context integration in language models.

---

## 4. Understanding mode connectivity via parameter space symmetry

URL: [View paper](#)

### Brief Assessment

Mode Connectivity Symmetry[22] focuses on parameter space symmetries and mode connectivity in neural networks, not on compositional adapters or context-to-parameter mappings. The theoretical framework addresses different problems: symmetry-induced connectivity versus compositional adapter generation.

---

## 5. Approximation of compositional functions with ReLU neural networks

URL: [View paper](#)

### Brief Assessment

Compositional ReLU Approximation[26] focuses on approximation theory for compositional functions using ReLU networks in control systems, not on parameter-space composition of adapters or context concatenation in language models.

---

## 6. On the realization of compositionality in neural networks

URL: [View paper](#)

### Brief Assessment

Compositionality Neural Networks[30] focuses on compositional understanding in neural networks through specialized neuron groups and modular structure, not on theoretical conditions for parameter-space composition via monoid homomorphisms or compositionality bounds for adapter addition.

---

## 7. Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial $\hat{\alpha}$

URL: [View paper](#)

### Brief Assessment

XPINNs[25] addresses domain decomposition for solving PDEs using neural networks with interface conditions between subdomains. This is fundamentally different from the original paper's focus on compositional adapters for language models and the theoretical framework for parameter-space composition in LLM adaptation.

---

## 8. Auto-GNN: Neural architecture search of graph neural networks

URL: [View paper](#)

### Brief Assessment

Auto-GNN[23] focuses on neural architecture search for graph neural networks, not on theoretical conditions for parameter-space composition in language models or adapter systems.

---

## 9. Finding Symmetry in Neural Network Parameter Spaces

URL: [View paper](#)

### Brief Assessment

Symmetry Parameter Spaces[29] focuses on loss-invariant transformations and parameter space symmetries in neural networks, not on adapter composition or context parametrization. The theoretical framework addresses different problems: [29] studies symmetry groups and their actions on parameters, while the original paper studies compositionality of adapters for context integration.

---

## 10. Skill Expansion and Composition in Parameter Space

URL: [View paper](#)

### Brief Assessment

Skill Expansion Composition[24] focuses on LoRA-based skill composition in reinforcement learning with diffusion policies, while the original paper addresses adapter composition for language models with context parametrization. The theoretical frameworks serve different purposes: [24] analyzes skill interpolation in RL parameter space, whereas the original establishes compositionality bounds for context-to-adapter mappings in LLMs.

---

## Contribution 3: Reversible context encoding with reconstruction capability

**Description:** The framework includes a reconstruction objective that allows the model to decode and recover the original input context from adapter parameters, providing a mechanism for verifying what information has been encoded and supporting safety and security requirements.

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. Bidirectional representations augmented autoregressive biological sequence generation: Application in de novo peptide sequencing

URL: [View paper](#)

### Brief Assessment

Bidirectional Autoregressive Generation[15] focuses on biological sequence generation (peptide sequencing) using a reconstruction objective to recover amino acid sequences from spectral data, not on reversible encoding of general context into model parameters for LLM adaptation as in the original paper.

---

## 2. Design and Optimization of Reversible Information Hiding Image Encryption Algorithms in the Context of Electronic Information Security

URL: [View paper](#)

### Brief Assessment

Reversible Information Hiding[13] focuses on image encryption with reversible data hiding in visual media, not on encoding context into model parameters with reconstruction capabilities for language models.

---

## 3. Motiondiffuser: Controllable multi-agent motion prediction using diffusion

URL: [View paper](#)

### Brief Assessment

Motiondiffuser[11] focuses on motion prediction for autonomous vehicles using diffusion models, not on encoding context into model parameters with reconstruction capabilities for language models.

---

## 4. Confidence Based Bidirectional Global Context Aware Training Framework for Neural Machine Translation

URL: [View paper](#)

### Brief Assessment

Confidence Bidirectional Training[20] uses reconstruction for training stability and faithfulness verification in NMT, not for reversible encoding of context into adapter parameters as in the original paper's compositional adapter framework.

---

## 5. Context-aware code generation with synchronous bidirectional decoder

URL: [View paper](#)

### Brief Assessment

Context Aware Generation[16] full text is not available (marked as 'n/a'), making comparison impossible. Without access to the candidate paper's content, no assessment of novelty refutation can be made.

---

## 6. A High-Capacity Reversible Data Hiding Scheme for Encrypted Hyperspectral Images Using Multi-Layer MSB Block Labeling and ERLE Compression

URL: [View paper](#)

### Brief Assessment

Reversible Data Hiding[12] focuses on reversible data hiding in encrypted hyperspectral images for secure data transmission, not on encoding context into model parameters with reconstruction capabilities for language models.

---

## 7. Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis

URL: [View paper](#)

### Brief Assessment

Imagebart[18] uses reconstruction for training a discrete autoencoder and verifying information in diffusion stages, not for encoding context into adapter parameters. The reconstruction serves a different purpose in a diffusion-based image synthesis framework rather than context parametrization for language models.

---

## 8. PointLIE: Locally Invertible Embedding for Point Cloud Sampling and Recovery

URL: [View paper](#)

### Brief Assessment

PointLIE[17] focuses on point cloud sampling and recovery using invertible networks to encode geometric offsets, not on encoding textual context into language model adapter parameters with reconstruction objectives for safety verification.

---

## 9. Memory Tokens: Large Language Models Can Generate Reversible Sentence Embeddings

URL: [View paper](#)

### Brief Assessment

Memory Tokens[19] focuses on generating reversible sentence embeddings for individual text sequences using optimized token embeddings, not on encoding multiple contexts into compositional adapter parameters. The reconstruction in Memory Tokens[19] serves to demonstrate perfect text recovery from a single embedding, whereas the original paper's reconstruction objective validates information preservation in composed adapters for multi-context integration.

---

## 10. RMem: Bridging Memory Retention and Retrieval via Reversible Compression

URL: [View paper](#)

### Prior Art Analysis

RMem[14] demonstrates prior work that implements reversible encoding of context into model parameters with reconstruction capability. Both papers employ a reversible architecture where context is compressed into model parameters and can be reconstructed. RMem[14] explicitly describes a 'reversible architecture, reconstructing raw data by invoking the model backward with compressed information' and uses 'bidirectional training with cycle consistency' to ensure faithful reconstruction. The original paper's claim to novelty in providing 'a reconstruction objective that allows the model to decode and recover the original input context from adapter parameters' is directly challenged by RMem[14]'s earlier implementation of the same concept.

### Evidence

Evidence 1 - **Rationale:** Both papers describe encoding information in a reversible manner with reconstruction capability. RMem[14] presents this as a core contribution, challenging the novelty claim of the original paper. - **Original:** furthermore, c ompas encodes information into adapter parameters in a reversible manner, enabling recovery of input context through a decoder, facilitating safety and security. - **Candidate:** we propose r3mem, a novel memory-augmented model that optimizes both memory retention and retrieval while minimizing external storage dependency. r3mem leverages a reversible architecture that integrates context compression and expansion, enabling assimilation and reconstruction of input data.

Evidence 2 - **Rationale:** Both papers use adapter-based approaches to encode context into parameters with reconstruction capability. RMem[14]'s bidirectional training with cycle consistency directly implements the reconstruction objective claimed as novel in the original paper. - **Original:** we introduce c ompas, a teacher-student framework for encoding context into adapters which facilitates composition in parameter space. (2) we establish theoretical conditions under which parameter-space addition provably approximates the behavior of the teacher model. (3) we empirically show that co... - **Candidate:** for retrieval, r3mem adopts a reversible architecture, reconstructing raw input by inverting the model invocation on compressed representations. this is achieved through adapter tuning, allowing seamless integration with pre-trained transformer model while maintaining parameter efficiency. to optimi...

Evidence 3 - **Rationale:** Both papers implement reconstruction of original context from compressed parameters. RMem[14]'s backward process that reconstructs raw content demonstrates the same reconstruction capability claimed in the original paper. - **Original:** to encourage faithfulness and improve training stability, we require adapters to recoverably encode the information contained in their contexts. we introduce a special query token [recon], which prompts the student to reconstruct the support context from its adapter in an autoregressive manner. - **Candidate:** the forward process encodes context into compressed memory representations, while the backward process reconstructs the raw content from memory tokens, enforcing consistency between the original and reconstructed information.

---

## Appendix: Text Similarity Detection

---

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

## References

---

- [0] Context Parametrization with Compositional Adapters [View paper](#)
- [1] Structural Priors and Modular Adapters in the Composable Fine-Tuning Algorithm of Large-Scale Models [View paper](#)
- [2] Corda: Context-oriented decomposition adaptation of large language models for task-aware parameter-efficient fine-tuning [View paper](#)
- [3] CorDA: Context-Oriented Decomposition Adaptation of Large Language Models [View paper](#)
- [4] Adapterfusion: Non-destructive task composition for transfer learning [View paper](#)
- [5] Generative adapter: Contextualizing language models in parameters with a single forward pass [View paper](#)
- [6] Latent symbol lattices in probabilistic semiosis: An unconventional architectural mechanism for contextual modulation in large language models [View paper](#)
- [7] Caila: concept-aware intra-layer adapters for compositional zero-shot learning [View paper](#)
- [8] On conditional and compositional language model differentiable prompting [View paper](#)
- [9] AdapterHub for Modular Natural Language Processing: The Complete Guide for Developers and Engineers [View paper](#)
- [10] A Context Awareness Framework for Real-Time Analysis and Adaptive Navigation in Dynamic Indoor Environments [View paper](#)
- [11] Motiondiffuser: Controllable multi-agent motion prediction using diffusion [View paper](#)
- [12] A High-Capacity Reversible Data Hiding Scheme for Encrypted Hyperspectral Images Using Multi-Layer MSB Block Labeling and ERLE Compression [View paper](#)
- [13] Design and Optimization of Reversible Information Hiding Image Encryption Algorithms in the Context of Electronic Information Security [View paper](#)
- [14] RMem: Bridging Memory Retention and Retrieval via Reversible Compression [View paper](#)
- [15] Bidirectional representations augmented autoregressive biological sequence generation: Application in de novo peptide sequencing [View paper](#)
- [16] Context-aware code generation with synchronous bidirectional decoder [View paper](#)
- [17] PointLIE: Locally Invertible Embedding for Point Cloud Sampling and Recovery [View paper](#)
- [18] Imagebart: Bidirectional context with multinomial diffusion for autoregressive image synthesis [View paper](#)
- [19] Memory Tokens: Large Language Models Can Generate Reversible Sentence Embeddings [View paper](#)
- [20] Confidence Based Bidirectional Global Context Aware Training Framework for Neural Machine Translation [View paper](#)
- [21] Theory of overparametrization in quantum neural networks [View paper](#)
- [22] Understanding mode connectivity via parameter space symmetry [View paper](#)
- [23] Auto-GNN: Neural architecture search of graph neural networks [View paper](#)
- [24] Skill Expansion and Composition in Parameter Space [View paper](#)
- [25] Extended physics-informed neural networks (XPINNs): A generalized space-time domain decomposition based deep learning framework for nonlinear partial  $\hat{\Delta}$  [View paper](#)
- [26] Approximation of compositional functions with ReLU neural networks [View paper](#)
- [27] Theoretical Investigation of Composite Neural Network [View paper](#)
- [28] Barron Spaces and the Compositional Function Spaces for Neural Network Models [View paper](#)
- [29] Finding Symmetry in Neural Network Parameter Spaces [View paper](#)
- [30] On the realization of compositionality in neural networks [View paper](#)
- [31] Human-like systematic generalization through a meta-learning neural network [View paper](#)
- [32] Neural-Fly enables rapid learning for agile flight in strong winds [View paper](#)
- [33] Meta-LoRA: Meta-Learning LoRA Components for Domain-Aware ID Personalization [View paper](#)
- [34] AdaFML: Adaptive Federated Meta Learning With Multi-Objectives and Context-Awareness in Dynamic Heterogeneous Networks [View paper](#)
- [35] TAML-Adapter: Enhancing Adapter Tuning Through Task-Agnostic Meta-Learning for Low-Resource Automatic Speech Recognition [View paper](#)
- [36] Adaptive compositional continual meta-learning [View paper](#)
- [37] Meta-learning hyperparameters for foundation model adaptation in remote-sensing imagery [View paper](#)
- [38] MDANet: A multi-stage domain adaptation framework for generalizable low-light image enhancement [View paper](#)
- [39] Neural Relational Inference with Fast Modular Meta-learning [View paper](#)
- [40] Modular meta-learning [View paper](#)