Novelty Assessment Report

Paper: Automated Stateful Specialization for Adaptive Agent Systems
PDF URL: https://openreview.net/pdf?id=UESTP6dR1K
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Current automated agent design frameworks produce either static workflows that lack adaptability or per-query optimizers that prevent the accumulation of deep, agent-level task expertise. We propose a new direction that reconciles these paradigms: creating stateful teams of specialist agents that accumulate knowledge over time and can be reconfigured for novel tasks entirely without human intervention. To this end, we introduce \textsc{ASpec}, a framework that manages this full agent lifecycle by first autonomously $\textbf{discovering}$ specialist archetypes via evolutionary search and then $\textbf{cultivating}$ their expertise through experience, mirroring how human experts learn through practice and reflection. We further introduce a lightweight hierarchical control policy, "retain-then-escalate," which governs when to leverage the established agent system versus when to adapt its structure. Through comprehensive experiments, we demonstrate that this approach leads to significant performance gains on expert-level scientific benchmarks like GPQA while matching the state-of-the-art on broader domain tasks, demonstrating a promising path toward agent systems that are simultaneously expert, adaptive, and efficient.

Core Task Landscape

This paper addresses: **Automated Discovery and Cultivation of Stateful Specialist Agents**
A total of **50 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

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

- **Agent Architecture and Memory Systems**
- **Multi-Agent Coordination and Specialization**
- **Task Planning and Execution**
- **Domain-Specific Applications**
- **Human-Computer Interaction and Interface Agents**
- **Machine Learning Support and Explainability**

Complete Taxonomy Tree

- Automated Discovery and Cultivation of Stateful Specialist Agents Survey Taxonomy
- Agent Architecture and Memory Systems
  - Stateful Agent Frameworks ★ (5 papers)
  - [0] Automated Stateful Specialization for Adaptive Agent Systems (Anon et al., 2026) View paper
  - [4] Agent-SAMA: State-Aware Mobile Assistant (Author(s), 2025) View paper
  - [8] State and Memory is All You Need for Robust and Reliable AI Agents (Matthew Muhoberac, 2025) View paper
  - [15] ALAS: A Stateful Multi-LLM Agent Framework for Disruption-Aware Planning (Chang, 2025) View paper
  - [23] Stateful active facilitator: Coordination and environmental heterogeneity in cooperative multi-agent reinforcement learning (Liu, 2022) View paper
  - Memory Models and Knowledge Retention (6 papers)
  - [11] Long term memory: The foundation of ai self-evolution (Jiang, 2024) View paper
  - [29] Adaptive Memory for Autonomous Agents (Stefan Tsokov, 2025) View paper
  - [34] Reflection-Based Memory For Web navigation Agents (Vempaty, 2025) View paper
  - [36] Semantic memory modeling and memory interaction in learning agents (Wenwen Wang, 2016) View paper
  - [38] Evolution in Simulation: AI-Agent School with Dual Memory for High-Fidelity Educational Dynamics (Sheng Jin, 2025) View paper
  - [43] Livia: An Emotion-Aware AR Companion Powered by Modular AI Agents and Progressive Memory Compression (Xi Rui, 2025) View paper
  - Context-Aware Reasoning and Adaptation (6 papers)
  - [21] ContextMate: a context-aware smart agent for efficient data analysis (Aamir Khan Jadoon, 2024) View paper
  - [24] Cooking With Agents: Designing Context-aware Voice Interaction (Razan Jaber, 2024) View paper
  - [25] Fast and efficient context-aware services (Danny Raz, 2006) View paper
  - [26] Narrative Memory based Intelligent Agents (Muhammad Bilal, 2025) View paper
  - [32] A context-aware approach to automated negotiation using reinforcement learning (Dan E. KrÃ¶hling, 2021) View paper
  - [39] Learning and adaptation of strategies in automated negotiations between context-aware agents (D. KrÃ¶hling, 2024) View paper
- Multi-Agent Coordination and Specialization
  - Hierarchical Multi-Agent Systems (6 papers)
  - [3] LangChain & LangGraph in Production: Architectures for Multi-Agent LLM Systems (Pelluru, 2025) View paper

- [5] Polaris: A safety-focused llm constellation architecture for healthcare (Mukherjee, 2024) View paper
- [9] HM-RAG: Hierarchical Multi-Agent Multimodal Retrieval Augmented Generation (Liu Pei, 2025) View paper
- [31] Multi-Agent Collaborative Framework for Intelligent IT Operations: An AOI System with Context-Aware Compression and Dynamic Task Scheduling (Zishan Bai, 2025) View paper
- [42] Reasoning-Aware Prompt Orchestration: A Foundation Model for Multi-Agent Language Model Coordination (Dhrif, 2025) View paper
- [48] S-AI: A Sparse Artificial Intelligence System Orchestrated by a Hormonal MetaAgent and Context-Aware Specialized Agents (SAID SLAOUI, n.d.) View paper
- Collaborative Agent Networks (5 papers)
- [7] Operational AGI: A Language-Based Approach for Adaptive Multi-Agent Systems (Marozau, 2025) View paper
- [19] CARA: A Hybrid Framework Integrating Swarm AI Agents and Knowledge Graphs for Advanced LLM Reasoning (Nirodya Pussadeniya, 2024) View paper
- [20] Automatic agent chaining for multimodal task support (R Manuvinakurike, 2025) View paper
- [45] AgentFlow: A Context Aware Multi-Agent Framework for Dynamic Agent Collaboration (Gayathri Nettem, 2025) View paper
- [46] MultiFuzz: A Dense Retrieval-based Multi-Agent System for Network Protocol Fuzzing (Wael, 2025) View paper
- Dynamic Role Assignment and Specialization (3 papers)
- [17] Trajectory-Class-Aware Multi-Agent Reinforcement Learning (Hyungho Na, 2025) View paper
- [28] ROIS: Role-Based Multi-Agent Collaboration by Context-Time-Aware Information Sharing (H Qi, 2025) View paper
- [37] SANet: A Semantic-aware Agentic AI Networking Framework for Cross-layer Optimization in 6G (Yong Xiao, 2025) View paper

- Task Planning and Execution
  - Dynamic Task Decomposition (2 papers)
  - [2] The hitchhiker's guide to autonomous research: A survey of scientific agents (Xin-Ming Wang, 2025) View paper
  - [14] Advancing Agentic Systems: Dynamic Task Decomposition, Tool Integration and Evaluation using Novel Metrics and Dataset (Ahmad, 2024) View paper
  - Tool Integration and Selection (1 papers)
  - [10] Towards Human-Guided, Data-Centric LLM Co-Pilots (Saveliev, 2025) View paper
  - Robustness and Error Recovery (1 papers)
  - [50] ReInAgent: A Context-Aware GUI Agent Enabling Human-in-the-Loop Mobile Task Navigation (Jia Hai-tao, 2025) View paper
- Domain-Specific Applications
  - Language and Translation Systems (2 papers)
  - [1] Preserving Cultural Identity with Context-Aware Translation Through Multi-Agent AI Systems (Rahman Abdur, 2025) View paper
  - [16] MAATS: A Multi-Agent Automated Translation System Based on MQM Evaluation (Wang, 2025) View paper
  - Software Engineering and Code Analysis (2 papers)
  - [12] RepoTransAgent: Multi-Agent LLM Framework for Repository-Aware Code Translation (Yin Xin, 2025) View paper
  - [49] RefAgent: A Multi-agent LLM-based Framework for Automatic Software Refactoring (Lamothe, 2025) View paper
  - Cybersecurity and Anomaly Detection (2 papers)
  - [13] MultiPhishGuard: An LLM-based Multi-Agent System for Phishing Email Detection (Spero Eric, 2025) View paper
  - [33] Multi-Agent Visual Reasoning for Out-of-Distribution Detection in Complex Road Environments (Jeonghyo Song, 2025) View paper
  - Network and Infrastructure Management (1 papers)
  - [18] Towards Specialized Wireless Networks Using an ML-Driven Radio Interface (Wojnar, 2025) View paper
  - Optimization and Resource Allocation (2 papers)
  - [6] A personalized automated bidding framework for fairness-aware online advertising (Haoqi Zhang, 2023) View paper
  - [35] Industrial symbiosis: Context-aware strategies for automated negotiation of smart contracts in peer-to-peer markets of prosumers (Dan E. Krohling, 2020) View paper
  - Data Analysis and Scientific Research (1 papers)
  - [41] Integrating Data-Driven Insights with Domain Expertise Using Agentic Conversational Analytics for Well Completions Optimization (C. J. S. Santiago, 2025) View paper
  - Education and Training (1 papers)
  - [44] Adaptive and personalized educational ubiquitous multi-agent system using context-awareness services and mobile devices (Oscar M. Salazar, 2015) View paper
  - Personal Finance and Privacy (1 papers)
  - [30] ASTRAFIN:- AI Financial Agent (Er. Jagpreet Singh, 2025) View paper
- Human-Computer Interaction and Interface Agents
  - GUI and Mobile Interface Agents (1 papers)
  - [47] MEGA-GUI: Multi-stage Enhanced Grounding Agents for GUI Elements (SeokJoo Kwak, 2025) View paper
- Machine Learning Support and Explainability
  - XAI and Interpretability (1 papers)
  - [40] Enhancing ML Explainability with Multi-Agent LLMs: A Context-Aware XAI Approach (R. O. Miyaji, 2025) View paper
  - Specialized Analysis and Rule Extraction (2 papers)
  - [22] Rcaa: Relational context-aware agents for person search (Xiaojun Chang, 2018) View paper
  - [27] Automatically Finding Rule-Based Neurons in OthelloGPT (Singh Aditya, 2025) View paper

Narrative

Core task: automated discovery and cultivation of stateful specialist agents. The field organizes around six main branches that reflect different facets of building and deploying intelligent agents. Agent Architecture and Memory Systems explores how agents maintain and leverage internal state over time, encompassing frameworks that integrate long-term memory, context-aware reasoning, and adaptive recall mechanisms. Multi-Agent Coordination and Specialization examines how collections of agents divide labor, negotiate roles, and collaborate on complex objectives. Task Planning and Execution focuses on decomposing goals into executable steps and orchestrating workflows, while Domain-Specific Applications tailors agent designs to verticals such as healthcare, finance, and cybersecurity. Human-Computer Interaction and Interface Agents address user-facing systems that guide, explain, or co-create with people, and Machine Learning Support and Explainability investigates how agents can interpret model behavior, provide transparency, and support decision-making. Together, these branches capture the spectrum from low-level memory primitives to high-level coordination and user engagement.

A particularly active line of work centers on stateful frameworks that enable agents to accumulate experience and refine their behavior across episodes. For instance, State Memory Agents[8] and ALAS[15] emphasize persistent memory structures that allow agents to recall past interactions and adapt strategies dynamically, while LangChain Production[3] and Agent SAMA[4] illustrate how production-ready architectures balance modularity with performance. Stateful Specialization[0] sits squarely within this cluster, focusing on mechanisms that not only preserve state but also guide the emergence of specialized roles through iterative learning. Compared to neighbors like Stateful Active Facilitator[23], which targets collaborative facilitation with explicit state management, Stateful Specialization[0] places greater emphasis on the automated discovery process itself—how agents identify niches and cultivate domain-specific expertise without exhaustive manual design. This distinction highlights an ongoing tension between hand-crafted specialization and emergent role differentiation, a theme that resonates across many branches as researchers seek scalable pathways to capable, context-aware agents.

Related Works in Same Category

The following **4 sibling papers** share the same taxonomy leaf node with the original paper:

1. Agent-SAMA: State-Aware Mobile Assistant

Authors: Anonymous Author(s) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Mobile Graphical User Interface (GUI) agents aim to autonomously complete tasks within or across apps based on user instructions. While recent Multimodal Large Language Models (MLLMs) enable these agents to interpret UI screens and perform actions, existing agents remain fundamentally reactive. They reason over the current UI screen but lack a structured representation of the app navigation flow, limiting GUI agents'ability to understand execution context, detect unexpected execution results, an...

Relationship Analysis

Both papers belong to the Stateful Agent Frameworks category, focusing on architectures that maintain persistent state and context across interactions. Agent-SAMA overlaps with the original paper in maintaining agent state through memory mechanisms (FSM-based state tracking vs. experience-driven memory), but differs fundamentally in scope: Agent-SAMA focuses specifically on mobile GUI navigation using Finite State Machines for execution verification and error recovery, while the original paper (ASPEC) addresses automated discovery and cultivation of specialist agents across diverse reasoning domains through evolutionary search and experience accumulation.

---

2. State and Memory is All You Need for Robust and Reliable AI Agents

Authors: Matthew Muhoberac, Atharva Parikh, Nirvi Vakharia, S. Virani, Aco Radujevic, et al. (19 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Large language models (LLMs) have enabled powerful advances in natural language understanding and generation. Yet their application to complex, real-world scientific workflows remain limited by challenges in memory, planning, and tool integration. Here, we introduce SciBORG (Scientific Bespoke Artificial Intelligence Agents Optimized for Research Goals), a modular agentic framework that allows LLM-based agents to autonomously plan, reason, and achieve robust and reliable domain-specific task exe...

Relationship Analysis

Both papers belong to the Stateful Agent Frameworks category, focusing on architectures that maintain persistent state and context across interactions. They overlap in their emphasis on memory systems for agent adaptation: the original paper uses evolutionary discovery and cultivation of specialist agents with episodic memory, while the candidate paper employs finite-state automaton (FSA) memory for scientific laboratory automation. The key difference is that the original paper focuses on automated discovery and cultivation of diverse specialist agents for general reasoning tasks, whereas the candidate paper targets robust execution of scientific workflows through FSA-based state tracking and hardware integration.

---

3. ALAS: A Stateful Multi-LLM Agent Framework for Disruption-Aware Planning

Authors: Chang, Edward Y., Edward Y. Chang, Longling Geng | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Large language models (LLMs) excel at rapid generation of text and multimodal content, yet they falter on transaction-style planning that demands ACID-like guarantees and real-time disruption recovery. We present Adaptive LLM Agent System (ALAS), a framework that tackles four fundamental LLM deficits: (i) absence of self-verification, (ii) context erosion, (iii) next-token myopia, and (iv) lack of persistent state. ALAS decomposes each plan into role-specialized agents, equips them with automati...

Relationship Analysis

Both papers belong to the Stateful Agent Frameworks category, focusing on architectures that maintain persistent state and context across interactions. They overlap in their emphasis on stateful memory systems and adaptive agent behavior, with both addressing the challenge of maintaining consistency across multi-step tasks. However, the original paper (ASPEC) focuses on evolutionary discovery and cultivation of specialist agents through automated search and experience accumulation for general task domains, while the candidate paper (ALAS) focuses on transaction-style planning with ACID-like guarantees and disruption recovery in specific domains like scheduling and logistics, using role-specialized agents with compensation protocols rather than evolutionary discovery.

---

4. Stateful active facilitator: Coordination and environmental heterogeneity in cooperative multi-agent reinforcement learning

Authors: Liu, Dianbo, Shah, Vedant, Dianbo Liu, et al. (27 authors total) | **Year/Venue**: 2022 | **URL**: View paper

Abstract

In cooperative multi-agent reinforcement learning, a team of agents works together to achieve a common goal. Different environments or tasks may require varying degrees of coordination among agents in order to achieve the goal in an optimal way. The nature of coordination will depend on the properties of the environment -- its spatial layout, distribution of obstacles, dynamics, etc. We term this variation of properties within an environment as heterogeneity. Existing literature has not sufficie...

Relationship Analysis

Both papers belong to the Stateful Agent Frameworks category, focusing on architectures that maintain persistent state and context across interactions. The original paper (ASPEC) addresses automated discovery and cultivation of specialist agents through evolutionary search and experience-based memory accumulation for task expertise, while the candidate paper (SAF) focuses on multi-agent coordination in heterogeneous environments using a shared knowledge source and dynamic policy selection from a pool. The key difference is that ASPEC emphasizes autonomous agent lifecycle management (discovery and cultivation of specialists) for single-task expertise, whereas SAF addresses real-time coordination challenges in cooperative multi-agent systems with varying environmental conditions.

Contributions Analysis

Overall novelty summary. The paper introduces ASpec, a framework for autonomously discovering and cultivating stateful specialist agents through evolutionary search and experience-based learning. It resides in the 'Stateful Agent Frameworks' leaf, which contains five papers total, indicating a moderately populated research direction within the broader Agent Architecture and Memory Systems branch. This leaf focuses specifically on architectures that maintain persistent state across interactions, distinguishing it from adjacent leaves that emphasize memory structures alone or context-aware reasoning without explicit state tracking.

The taxonomy reveals that ASpec sits at the intersection of multiple research threads. Its nearest neighbors include Memory Models and Knowledge Retention (six papers on long-term storage and knowledge evolution) and Context-Aware Reasoning and Adaptation (six papers on dynamic behavioral adjustment). The Multi-Agent Coordination and Specialization branch, particularly Dynamic Role Assignment and Specialization (three papers), addresses related themes of role learning but typically through reinforcement or fixed architectures rather than evolutionary discovery. The scope notes clarify that ASpec's lifecycle management approach—spanning discovery, cultivation, and hierarchical control—bridges state management with coordination mechanisms that traditionally belong to separate branches.

Among 26 candidates examined across three contributions, none were flagged as clearly refuting the work. The ASpec framework contribution examined 10 candidates with zero refutable overlaps, the retain-then-escalate policy examined 6 with none refutable, and the paradigm reconciliation examined 10 with none refutable. This suggests that within the limited search scope, the combination of evolutionary archetype discovery, experience-based cultivation, and lightweight hierarchical control appears distinct from prior stateful frameworks like State Memory Agents or ALAS, which emphasize memory persistence but not autonomous role emergence. The absence of refutable candidates indicates novelty relative to the examined set, though the search scale leaves open the possibility of relevant work beyond the top-26 semantic matches.

Based on the limited literature search, ASpec appears to occupy a relatively sparse niche within stateful agent research, particularly in its emphasis on autonomous lifecycle management without human intervention. The taxonomy structure shows that while stateful frameworks and multi-agent specialization are active areas, the specific integration of evolutionary discovery with hierarchical control policies is less densely explored. However, the analysis covers only top-26 semantic matches and does not exhaustively survey adjacent fields like reinforcement learning for role assignment or meta-learning for agent adaptation, which may contain relevant prior work not captured in this scope.

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

Contribution 1: ASPEC framework for stateful specialist agent lifecycle management

Description: ASPEC is a framework that automates the complete lifecycle of specialist agents through two phases: evolutionary discovery of agent archetypes and autonomous cultivation of their expertise through experience accumulation. This enables agents to develop deep, persistent knowledge over time rather than being regenerated for each query.

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. Autonomous agents and multi-agent systems: explorations in learning, self-organization and adaptive computation

URL: View paper

Brief Assessment

Learning Self Organization[68] appears to focus on evolutionary approaches to autonomous agents in general contexts (e.g., image processing), not on the specific two-phase lifecycle framework (discovery + cultivation) for stateful specialist agents with persistent memory that ASPEC proposes.

---

2. From Agentification to Self-Evolving Agentic AI for Wireless Networks: Concepts, Approaches, and Future Research Directions

URL: View paper

Brief Assessment

Agentification Wireless[63] focuses on self-evolving agents in wireless network contexts with hardware-software co-evolution (e.g., antenna optimization), not on automated lifecycle management of specialist agents through evolutionary discovery and expertise cultivation as in ASPEC.

---

3. The Evolution of Agentic AI: Architecture and Workflows for Autonomous Systems

URL: View paper

Brief Assessment

Agentic AI Evolution[69] provides a high-level overview of agentic architectures and workflows but does not describe any specific framework for automated lifecycle management, evolutionary discovery, or expertise cultivation of specialist agents.

---

4. Autonomous Industrial Management via Reinforcement Learning: Towards Self-Learning Agents for Decision-Making

URL: View paper

Brief Assessment

Autonomous Industrial Management[66] focuses on industrial management systems with reinforcement learning for decision-making in manufacturing contexts, not on automated lifecycle management of specialist agents through evolutionary discovery and expertise cultivation as described in ASPEC.

---

5. Orchestrating Autonomy: Patterns, Protocols, and Governance for Enterprise Agentic AI

URL: View paper

Brief Assessment

Orchestrating Autonomy[70] focuses on enterprise governance patterns and coordination protocols for agentic AI systems, not on automated lifecycle management frameworks with evolutionary discovery and autonomous expertise cultivation mechanisms.

---

6. Enabling Autonomic Microservice Management through Self-Learning Agents

URL: View paper

Brief Assessment

Autonomic Microservice Management[64] focuses on microservice management through curriculum learning and exploration tasks, not on general agent lifecycle management with evolutionary discovery and expertise cultivation as in ASPEC.

---

7. ATHENA: Agentic Team for Hierarchical Evolutionary Numerical Algorithms

URL: View paper

Brief Assessment

ATHENA[67] focuses on autonomous computational research lifecycle management for scientific computing and machine learning, using a HENA loop framed as a contextual bandit problem. This differs fundamentally from ASPEC's evolutionary discovery and cultivation of stateful specialist agents with persistent memory for general task domains.

---

8. Modeling adaptive autonomous agents

URL: View paper

Brief Assessment

Adaptive Autonomous Agents[65] focuses on modeling biological organisms as programs for evolutionary and behavioral studies, not on automated lifecycle management frameworks for specialist agents with evolutionary discovery and expertise cultivation phases.

---

9. AgentEvolver: Towards Efficient Self-Evolving Agent System

URL: View paper

Brief Assessment

AgentEvolver[61] focuses on self-evolving agents through curiosity-driven task generation, experience reuse, and differentiated reward attribution in RL pipelines. It does not address the lifecycle management of stateful specialist agents with evolutionary discovery and autonomous expertise cultivation as described in ASPEC.

---

10. A comprehensive survey of self-evolving ai agents: A new paradigm bridging foundation models and lifelong agentic systems

URL: View paper

Brief Assessment

Self Evolving Survey[62] provides a broad taxonomy of agent evolution techniques but does not present a specific framework for automated lifecycle management of stateful specialists through evolutionary discovery and autonomous cultivation phases as described in ASPEC.

Contribution 2: Retain-then-escalate hierarchical control policy

Description: A lightweight hierarchical control policy that balances efficiency and adaptability by maintaining persistent specialist teams across queries and only triggering architectural redesign when necessary. This approach avoids the rediscovery costs of per-query regeneration while preserving adaptability.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. A Data Acquisition Method Based Collaborative Information Model for Heterogeneous Terminal Device Groups in Converter Stations

URL: View paper

Brief Assessment

Heterogeneous Terminal Groups[73] focuses on data acquisition in converter stations using multi-agent collaboration for sampling frequency and compression ratio adjustments. It does not address hierarchical control policies for balancing persistent specialist agent teams with adaptive architectural resampling in general agent systems.

---

2. RCTAMP: Enhancing Rule-Constrained TAMP via Multi-agent Closed-Loop Collaboration Integrating Consensus Planning

URL: View paper

Brief Assessment

RCTAMP[72] focuses on task and motion planning with replanning/resampling triggered by consistency checks, not on balancing persistent specialist teams with adaptive architectural redesign in multi-agent systems.

---

3. A multi-layered AI-driven cybersecurity architecture: Integrating entropy analytics, Fuzzy reasoning, game theory and multi-agent reinforcement learning for adaptive â□¦

URL: View paper

Brief Assessment

Multi Layered Cybersecurity[71] focuses on cybersecurity architecture with multi-agent reinforcement learning for security solutions, not on hierarchical control policies for balancing persistent agent teams with adaptive architectural resampling in general agent systems.

---

4. Automated Specialization of Stateful Agent Systems

URL: View paper

Brief Assessment

Automated Stateful Specialization[76] presents the identical 'retain-then-escalate' policy with the same technical implementation. This is the same paper, not a refutation of novelty.

---

5. Method of synthesis and automatic adaptation of the architecture of a hierarchical multi-agent system

URL: View paper

Brief Assessment

Hierarchical Synthesis Adaptation[75] focuses on modifying agent connections and using genetic algorithms for parameter optimization in hierarchical MAS, not on balancing persistent specialist teams with adaptive architectural resampling through a learned control policy.

---

6. A Survey on Reliability, Transparency, Accountability, and Fairness in LLM-based Multi-Agent Systems through the Responsibility Lens

URL: View paper

Brief Assessment

Responsibility Lens Survey[74] is a survey paper on reliability, transparency, accountability, and fairness in LLM-based multi-agent systems. The provided context fragments mention hierarchical structures and adaptive graphs but do not describe a control policy mechanism comparable to the retain-then-escalate approach for balancing persistent specialist teams with architectural resampling.

Contribution 3: Reconciliation of task-level and query-level agent design paradigms

Description: The framework bridges the gap between static task-level optimization and dynamic query-level adaptation by enabling specialist agents to maintain persistent expertise while remaining adaptable to new tasks through automated reconfiguration without human intervention.

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. Building self-evolving agents via experience-driven lifelong learning: A framework and benchmark

URL: View paper

Brief Assessment

Self Evolving Agents[57] focuses on experience-driven lifelong learning through continuous environmental interaction and skill abstraction from trajectories, not on reconciling static task-level optimization with dynamic query-level adaptation in agent design automation.

---

2. El Agente: An autonomous agent for quantum chemistry

URL: View paper

Brief Assessment

El Agente[51] focuses on quantum chemistry workflows with hierarchical task decomposition for domain-specific calculations, not on reconciling task-level versus query-level agent design paradigms in general automated agent systems.

---

3. Knowledge Retention for Continual Model-Based Reinforcement Learning

URL: View paper

Brief Assessment

Knowledge Retention[52] focuses on continual model-based reinforcement learning for world models across tasks with different reward functions, not on reconciling static task-level optimization with dynamic query-level adaptation in multi-agent systems.

---

4. Dynamic Knowledge Management in an Agent-Based Extended Green Cloud Simulator

URL: View paper

Brief Assessment

Dynamic Knowledge Management[60] focuses on runtime adaptation of system knowledge in cloud infrastructure using rule-based expert systems, not on reconciling static task-level optimization with dynamic query-level adaptation in agent design frameworks.

---

5. How artificially intelligent conversational agents influence EFL learners' self-regulated learning and retention

URL: View paper

Brief Assessment

Conversational Agents Learning[55] focuses on AI conversational agents for English language learning and self-regulated learning, not on agent design frameworks or the reconciliation of task-level versus query-level optimization paradigms in automated agent systems.

---

6. Continual Knowledge Adaptation for Reinforcement Learning

URL: View paper

Brief Assessment

Continual Knowledge Adaptation[54] addresses continual reinforcement learning with task-specific knowledge vectors for non-stationary environments, not the reconciliation of static task-level optimization versus dynamic query-level adaptation in multi-agent system design.

---

7. SEAgent: Self-Evolving Computer Use Agent with Autonomous Learning from Experience

URL: View paper

Brief Assessment

SEAgent[58] focuses on autonomous evolution of computer-use agents through experiential learning in software environments, not on reconciling task-level versus query-level agent design paradigms or managing stateful specialist teams with automated reconfiguration.

---

8. Moral: Moe augmented lora for llms' lifelong learning

URL: View paper

Brief Assessment

Moral[59] focuses on lifelong learning for LLMs through mixture-of-experts augmented LoRA, addressing knowledge updating and catastrophic forgetting. This is fundamentally different from the original paper's contribution of reconciling static task-level optimization with dynamic query-level adaptation through stateful specialist agents with automated reconfiguration.

---

9. Improving LLM Agent Planning with In-Context Learning via Atomic Fact Augmentation and Lookahead Search

URL: View paper

Brief Assessment

Atomic Fact Augmentation[53] focuses on in-context learning through atomic fact extraction and lookahead search for single-agent planning, not on reconciling static task-level optimization with dynamic query-level adaptation through stateful specialist agent teams.

---

10. DRAE: Dynamic Retrieval-Augmented Expert Networks for Lifelong Learning and Task Adaptation in Robotics

URL: View paper

Brief Assessment

DRAE[56] focuses on lifelong learning in robotics through mixture-of-experts and retrieval-augmented generation, not on reconciling static task-level versus dynamic query-level agent design paradigms in multi-agent systems.

Appendix: Text Similarity Detection

Textual similarity detection checked 30 papers and found 4 similarity segment(s) across 2 paper(s).

The following **2 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. Automated Specialization of Stateful Agent Systems

Detected in: 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.

---

2. A comprehensive survey of self-evolving ai agents: A new paradigm bridging foundation models and lifelong agentic systems

Detected in: Contribution: contribution_1

⚠ **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] Automated Stateful Specialization for Adaptive Agent Systems View paper
- [1] Preserving Cultural Identity with Context-Aware Translation Through Multi-Agent AI Systems View paper
- [2] The hitchhiker's guide to autonomous research: A survey of scientific agents View paper
- [3] LangChain & LangGraph in Production: Architectures for Multi-Agent LLM Systems View paper
- [4] Agent-SAMA: State-Aware Mobile Assistant View paper
- [5] Polaris: A safety-focused llm constellation architecture for healthcare View paper
- [6] A personalized automated bidding framework for fairness-aware online advertising View paper
- [7] Operational AGI: A Language-Based Approach for Adaptive Multi-Agent Systems View paper
- [8] State and Memory is All You Need for Robust and Reliable AI Agents View paper
- [9] HM-RAG: Hierarchical Multi-Agent Multimodal Retrieval Augmented Generation View paper
- [10] Towards Human-Guided, Data-Centric LLM Co-Pilots View paper
- [11] Long term memory: The foundation of ai self-evolution View paper
- [12] RepoTransAgent: Multi-Agent LLM Framework for Repository-Aware Code Translation View paper
- [13] MultiPhishGuard: An LLM-based Multi-Agent System for Phishing Email Detection View paper
- [14] Advancing Agentic Systems: Dynamic Task Decomposition, Tool Integration and Evaluation using Novel Metrics and Dataset View paper
- [15] ALAS: A Stateful Multi-LLM Agent Framework for Disruption-Aware Planning View paper
- [16] MAATS: A Multi-Agent Automated Translation System Based on MQM Evaluation View paper
- [17] Trajectory-Class-Aware Multi-Agent Reinforcement Learning View paper
- [18] Towards Specialized Wireless Networks Using an ML-Driven Radio Interface View paper
- [19] CARA: A Hybrid Framework Integrating Swarm AI Agents and Knowledge Graphs for Advanced LLM Reasoning View paper
- [20] Automatic agent chaining for multimodal task support View paper
- [21] ContextMate: a context-aware smart agent for efficient data analysis View paper
- [22] Rcaa: Relational context-aware agents for person search View paper
- [23] Stateful active facilitator: Coordination and environmental heterogeneity in cooperative multi-agent reinforcement learning View paper
- [24] Cooking With Agents: Designing Context-aware Voice Interaction View paper
- [25] Fast and efficient context-aware services View paper
- [26] Narrative Memory based Intelligent Agents View paper
- [27] Automatically Finding Rule-Based Neurons in OthelloGPT View paper
- [28] ROIS: Role-Based Multi-Agent Collaboration by Context-Time-Aware Information Sharing View paper
- [29] Adaptive Memory for Autonomous Agents View paper
- [30] ASTRAFIN:- AI Financial Agent View paper
- [31] Multi-Agent Collaborative Framework for Intelligent IT Operations: An AOI System with Context-Aware Compression and Dynamic Task Scheduling View paper
- [32] A context-aware approach to automated negotiation using reinforcement learning View paper
- [33] Multi-Agent Visual Reasoning for Out-of-Distribution Detection in Complex Road Environments View paper
- [34] Reflection-Based Memory For Web navigation Agents View paper
- [35] Industrial symbiosis: Context-aware strategies for automated negotiation of smart contracts in peer-to-peer markets of prosumers View paper
- [36] Semantic memory modeling and memory interaction in learning agents View paper
- [37] SANet: A Semantic-aware Agentic AI Networking Framework for Cross-layer Optimization in 6G View paper
- [38] Evolution in Simulation: AI-Agent School with Dual Memory for High-Fidelity Educational Dynamics View paper
- [39] Learning and adaptation of strategies in automated negotiations between context-aware agents View paper
- [40] Enhancing ML Explainability with Multi-Agent LLMs: A Context-Aware XAI Approach View paper
- [41] Integrating Data-Driven Insights with Domain Expertise Using Agentic Conversational Analytics for Well Completions Optimization View paper
- [42] Reasoning-Aware Prompt Orchestration: A Foundation Model for Multi-Agent Language Model Coordination View paper
- [43] Livia: An Emotion-Aware AR Companion Powered by Modular AI Agents and Progressive Memory Compression View paper
- [44] Adaptive and personalized educational ubiquitous multi-agent system using context-awareness services and mobile devices View paper
- [45] AgentFlow: A Context Aware Multi-Agent Framework for Dynamic Agent Collaboration View paper
- [46] MultiFuzz: A Dense Retrieval-based Multi-Agent System for Network Protocol Fuzzing View paper
- [47] MEGA-GUI: Multi-stage Enhanced Grounding Agents for GUI Elements View paper
- [48] S-AI: A Sparse Artificial Intelligence System Orchestrated by a Hormonal MetaAgent and Context-Aware Specialized Agents View paper
- [49] RefAgent: A Multi-agent LLM-based Framework for Automatic Software Refactoring View paper
- [50] ReInAgent: A Context-Aware GUI Agent Enabling Human-in-the-Loop Mobile Task Navigation View paper

- [51] El Agente: An autonomous agent for quantum chemistry View paper
- [52] Knowledge Retention for Continual Model-Based Reinforcement Learning View paper
- [53] Improving LLM Agent Planning with In-Context Learning via Atomic Fact Augmentation and Lookahead Search View paper
- [54] Continual Knowledge Adaptation for Reinforcement Learning View paper
- [55] How artificially intelligent conversational agents influence EFL learners' self-regulated learning and retention View paper
- [56] DRAE: Dynamic Retrieval-Augmented Expert Networks for Lifelong Learning and Task Adaptation in Robotics View paper
- [57] Building self-evolving agents via experience-driven lifelong learning: A framework and benchmark View paper
- [58] SEAgent: Self-Evolving Computer Use Agent with Autonomous Learning from Experience View paper
- [59] Moral: Moe augmented lora for llms' lifelong learning View paper
- [60] Dynamic Knowledge Management in an Agent-Based Extended Green Cloud Simulator View paper
- [61] AgentEvolver: Towards Efficient Self-Evolving Agent System View paper
- [62] A comprehensive survey of self-evolving ai agents: A new paradigm bridging foundation models and lifelong agentic systems View paper
- [63] From Agentification to Self-Evolving Agentic AI for Wireless Networks: Concepts, Approaches, and Future Research Directions View paper
- [64] Enabling Autonomic Microservice Management through Self-Learning Agents View paper
- [65] Modeling adaptive autonomous agents View paper
- [66] Autonomous Industrial Management via Reinforcement Learning: Towards Self-Learning Agents for Decision-Making View paper
- [67] ATHENA: Agentic Team for Hierarchical Evolutionary Numerical Algorithms View paper
- [68] Autonomous agents and multi-agent systems: explorations in learning, self-organization and adaptive computation View paper
- [69] The Evolution of Agentic AI: Architecture and Workflows for Autonomous Systems View paper
- [70] Orchestrating Autonomy: Patterns, Protocols, and Governance for Enterprise Agentic AI View paper
- [71] A multi-layered AI-driven cybersecurity architecture: Integrating entropy analytics, Fuzzy reasoning, game theory and multi-agent reinforcement learning for adaptive â View paper
- [72] RCTAMP: Enhancing Rule-Constrained TAMP via Multi-agent Closed-Loop Collaboration Integrating Consensus Planning View paper
- [73] A Data Acquisition Method Based Collaborative Information Model for Heterogeneous Terminal Device Groups in Converter Stations View paper
- [74] A Survey on Reliability, Transparency, Accountability, and Fairness in LLM-based Multi-Agent Systems through the Responsibility Lens View paper
- [75] Method of synthesis and automatic adaptation of the architecture of a hierarchical multi-agent system View paper
- [76] Automated Specialization of Stateful Agent Systems View paper