

Novelty Assessment Report

Paper: KnowProxy: Adapting Large Language Models by Knowledge-guided Proxy

PDF URL: <https://openreview.net/pdf?id=14f18NoEqO>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Adapting large language models (LLMs) using smaller proxy models has been shown to improve training efficiency, where the LLMs remain frozen while the proxies are tuned on top. However, this approach typically requires access to the output probability distributions of LLMs, which are often inaccessible or unstable. To address this limitation, we propose KnowProxy, a knowledge-guided proxy framework in which the proxy is trained with textual knowledge rather than probability distributions. Specifically, we first elicit textual knowledge and reasoning from frozen LLMs through prompting, and then the proxy model learns to adapt this reasoning to target task distributions. We evaluate KnowProxy on diverse reasoning benchmarks with different fine-tuning scenarios. Comprehensive results show that KnowProxy achieves competitive or even better performance without direct access to probability distributions, thereby providing a scalable and versatile alternative to traditional fine-tuning.

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

Core Task Landscape

This paper addresses: **Adapting Large Language Models through Knowledge-Guided Proxy Training**

A total of **34 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

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

- **Proxy-Based Model Adaptation Mechanisms**
- **Knowledge-Guided Adaptation Strategies**
- **Parameter-Efficient Adaptation Strategies**
- **Knowledge Transfer and Distillation Methods**
- **Domain Adaptation and Transfer Learning**
- **Context and Knowledge Sensitivity Control**
- **Specialized Application Domains**

Complete Taxonomy Tree

- Adapting Large Language Models through Knowledge-Guided Proxy Training Survey Taxonomy
- Proxy-Based Model Adaptation Mechanisms
 - Decoding-Time Proxy Tuning (2 papers)
 - [6] Tuning Language Models by Proxy (Liu, 2024) [View paper](#)
 - [14] Improving instruction following in language models through proxy-based uncertainty estimation (Lee Joon-Ho, 2024) [View paper](#)
 - Training-Time Proxy Integration ★ (3 papers)
 - [0] KnowProxy: Adapting Large Language Models by Knowledge-guided Proxy (Anon et al., 2026) [View paper](#)
 - [21] FedPromo: Federated Lightweight Proxy Models at the Edge Bring New Domains to Foundation Models (Barbato, 2025) [View paper](#)
 - [23] Large and Small Model Collaboration for Air Interface (Yiming Cui, 2025) [View paper](#)
 - Proxy-Based Architecture Search and Optimization (2 papers)
 - [11] Lpzero: Language model zero-cost proxy search from zero (Dong Pei-jie, 2024) [View paper](#)
 - [15] Revolutionizing Mixed Precision Quantization: Towards Training-free Automatic Proxy Discovery via Large Language Models (Haidong Kang, 2025) [View paper](#)
- Knowledge-Guided Adaptation Strategies
 - Structured Knowledge Integration (3 papers)
 - [2] Bridging Stepwise Lab-Informed Pretraining and Knowledge-Guided Learning for Diagnostic Reasoning (Hu, 2024) [View paper](#)
 - [16] InfuserKI: Enhancing Large Language Models with Knowledge Graphs via Infuser-Guided Knowledge Integration (Wang FaLi, 2024) [View paper](#)
 - [31] A Multi-Modal Knowledge-Enhanced Framework for Vessel Trajectory Prediction (Haomin Yu, 2025) [View paper](#)
 - Domain-Specific Knowledge Adaptation (3 papers)
 - [3] KG-SR-LLM: Knowledge-Guided Semantic Representation and Large Language Model Framework for Cross-Domain Bearing Fault Diagnosis (Chengyong Xiao, 2025) [View paper](#)
 - [5] Knowledge and Task-Driven Multimodal Adaptive Transfer Through LLMs with Limited Data (Xu Zhang, 2024) [View paper](#)
 - [7] Adapting Large Language Models to Log Analysis with Interpretable Domain Knowledge (Yuhe Ji, 2024) [View paper](#)
 - Knowledge-Guided Pre-training and Verification (2 papers)
 - [25] Augmented Large Language Models with Parametric Knowledge Guiding (Luo, 2023) [View paper](#)
 - [26] KgPLM: Knowledge-guided Language Model Pre-training via Generative and Discriminative Learning (He Bin, 2022) [View paper](#)

- Retrieval-Augmented Adaptation (1 papers)
- [34] Self-Knowledge Guided Retrieval Augmentation for Large Language Models (Wang, 2023) [View paper](#)
- Parameter-Efficient Adaptation Strategies
 - Low-Rank and Singular-Value Adaptation (2 papers)
 - [1] KaSA: Knowledge-Aware Singular-Value Adaptation of Large Language Models (Wang Fan, 2024) [View paper](#)
 - [8] Enhancing Multimodal LLM for Detailed and Accurate Video Captioning using Multi-Round Preference Optimization (Tang, 2024) [View paper](#)
 - Modular Adapter Integration (2 papers)
 - [22] Selective Knowledge Injection via Adapter Modules in Large-Scale Language Models (Hongye Zheng, 2025) [View paper](#)
 - [30] NoEsis: Differentially Private Knowledge Transfer in Modular LLM Adaptation (Romijnders, 2025) [View paper](#)
 - Cross-Modal Prompt Learning (2 papers)
 - [4] KN-VLM: KNowledge-guided Vision-and-Language Model for visual abductive reasoning (Kuo Tan, 2025) [View paper](#)
 - [10] Hierarchical cross-modal prompt learning for vision-language models (Zheng Hao, 2025) [View paper](#)
- Knowledge Transfer and Distillation Methods
 - Preference-Based Knowledge Distillation (2 papers)
 - [12] Aligning Teacher with Student Preferences for Tailored Training Data Generation (Liu Yan-tao, 2024) [View paper](#)
 - [13] Direct Preference Knowledge Distillation for Large Language Models (Li, 2024) [View paper](#)
 - Bidirectional and Flipped Distillation (1 papers)
 - [17] Flipping Knowledge Distillation: Leveraging Small Models' Expertise to Enhance LLMs in Text Matching (Mingzhe Li, 2025) [View paper](#)
 - Federated and Distributed Knowledge Transfer (1 papers)
 - [24] FedLMA: A Federated Learning Framework Integrating LLM-Based Multi-Agent Reasoning With Knowledge Distillation (Jin-Long Wang, 2025) [View paper](#)
- Domain Adaptation and Transfer Learning
 - Source-Free Domain Adaptation (2 papers)
 - [29] Deciphering Invariant Feature Decoupling in Source-free Time Series Forecasting with Proxy Denoising (Liu Chenxi, 2025) [View paper](#)
 - [33] Proxy Denoising for Source-Free Domain Adaptation (Tang Song, 2024) [View paper](#)
 - Multi-Domain and Cross-Domain Transfer (1 papers)
 - [9] USTC-BUPT at SemEval-2024 Task 8: Enhancing Machine-Generated Text Detection via Domain Adversarial Neural Networks and LLM Embeddings (Zikang Guo, 2024) [View paper](#)
- Context and Knowledge Sensitivity Control
 - Dynamic Knowledge Sensitivity Steering (1 papers)
 - [18] Continuously Steering LLMs Sensitivity to Contextual Knowledge with Proxy Models (Wang Yi-lin, 2025) [View paper](#)
 - Knowledge Conflict and Unlearning Management (1 papers)
 - [28] Towards Mitigating Excessive Forgetting in LLM Unlearning via Entanglement-Aware Unlearning with Proxy Constraint (Liu Zhi-hao, 2025) [View paper](#)
- Specialized Application Domains
 - Multi-Agent and Embodied Environment Adaptation (1 papers)
 - [20] Learn as Individuals, Evolve as a Team: Multi-agent LLMs Adaptation in Embodied Environments (Li, 2025) [View paper](#)
 - Audio-Visual and Multimodal Segmentation (1 papers)
 - [19] OpenAVS: Training-Free Open-Vocabulary Audio Visual Segmentation with Foundational Models (Chen Sheng-Kai, 2025) [View paper](#)
 - Medical and Clinical Decision Support (2 papers)
 - [27] Large language model uncertainty proxies: discrimination and calibration for medical diagnosis and treatment. (Thomas Savage, 2024) [View paper](#)
 - [32] Incorporating Patient Values in Large Language Model Recommendations for Surrogate and Proxy Decisions. (Jeremy A Balch, 2024) [View paper](#)

Narrative

Core task: adapting large language models through knowledge-guided proxy training. This field addresses the challenge of efficiently adapting large language models by leveraging smaller proxy models or auxiliary knowledge sources during training. The taxonomy reveals several complementary branches: Proxy-Based Model Adaptation Mechanisms explore how smaller models can guide or substitute for expensive large-model updates, including training-time integration approaches like KnowProxy[0] and inference-time strategies such as Tuning by Proxy[6]. Knowledge-Guided Adaptation Strategies focus on injecting structured or domain-specific knowledge into models, with works like KaSA[1] and KG-SR-LLM[3] demonstrating how external knowledge graphs or task-driven cues can steer adaptation. Parameter-Efficient Adaptation Strategies and Knowledge Transfer and Distillation Methods address scalability through techniques like low-rank updates and student-teacher frameworks, while Domain Adaptation branches tackle specialized settings from medical imaging to federated learning, and Context and Knowledge Sensitivity Control manages how models balance parametric versus contextual information.

A particularly active line of work centers on training-time proxy integration, where smaller models serve as computational surrogates to guide large-model fine-tuning without full-scale backpropagation. KnowProxy[0] exemplifies this approach by using knowledge-guided proxies during training, sitting naturally alongside FedPromo[21] and Large Small Collaboration[23], which similarly exploit small-large model synergies in federated and collaborative settings. These methods contrast with knowledge injection strategies like Selective Knowledge Injection[22] or Parametric Knowledge Guiding[25], which emphasize embedding external structured knowledge rather than relying on proxy architectures. The trade-off revolves around whether adaptation should prioritize computational efficiency through architectural proxies or semantic richness through explicit knowledge integration. KnowProxy[0] bridges these themes by combining proxy-based efficiency with knowledge-guided steering, positioning it at the intersection of mechanistic innovation and knowledge-aware adaptation within the broader landscape of parameter-efficient LLM tuning.

Related Works in Same Category

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

1. FedPromo: Federated Lightweight Proxy Models at the Edge Bring New Domains to Foundation Models

Authors: Barbato, Francesco, Shenaj, Donald, Michieli, et al. (8 authors total) | **Year/Venue:** 2025 • arXiv.org | **URL:** [View paper](#)

Abstract

Federated Learning (FL) is an established paradigm for training deep learning models on decentralized data. However, as the size of the models grows, conventional FL approaches often require significant computational resources on client devices, which may not be feasible. We introduce FedPromo, a novel framework that enables efficient adaptation of large-scale foundation models stored on a central server to new domains encountered only by remote clients. Instead of directly training the large mo...

Relationship Analysis

Both papers belong to the Training-Time Proxy Integration category, where smaller proxy models are trained to adapt larger models during training phases. While KnowProxy focuses on adapting LLMs through textual knowledge and reasoning elicited via prompting (with the proxy learning to map this knowledge to task distributions), FedPromo addresses federated learning scenarios where lightweight proxy models are trained on client devices and aggregated to adapt foundation models on a central server. The key difference is that KnowProxy operates in a centralized setting using knowledge distillation from LLM-generated text, whereas FedPromo tackles decentralized, privacy-preserving adaptation across distributed clients with resource constraints.

2. Large and Small Model Collaboration for Air Interface

Authors: Yiming Cui, Jiajia Guo, Xiao Li, Chao-Kai Wen, Shi Jin | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

Large artificial intelligence models (LAMs) have shown strong capability in wireless communications, yet existing works mainly rely on their generalized knowledge across environments while overlooking the potential gains of environment-specific adaptation. Directly fine-tuning LAMs for adaptation is often impractical due to prohibitive training costs, low inference efficiency in multi-user scenarios, and the risk of catastrophic forgetting, in addition to the limited accessibility of model param...

Relationship Analysis

Both papers belong to the Training-Time Proxy Integration category, where smaller proxy models are trained to adapt larger models during training phases. They overlap in using proxy models to avoid direct fine-tuning of large models and in training proxies on knowledge derived from the larger model. However, KnowProxy focuses on adapting LLMs for NLP reasoning tasks using textual knowledge elicited through prompting, while the candidate paper (LASCO) addresses wireless communication CSI feedback tasks by training proxies to learn environment-specific reconstruction shifts from a base LAM, representing distinct application domains and knowledge transfer mechanisms.

Contributions Analysis

Overall novelty summary. The paper proposes KnowProxy, a framework that trains smaller proxy models using textual knowledge elicited from frozen large language models rather than requiring access to their probability distributions. Within the taxonomy, this work resides in the Training-Time Proxy Integration leaf under Proxy-Based Model Adaptation Mechanisms, alongside two sibling papers (FedPromo and Large Small Collaboration). This leaf represents a moderately populated research direction within a broader taxonomy of 34 papers across 20 leaf nodes, indicating focused but not overcrowded attention to training-time proxy strategies for LLM adaptation.

The taxonomy structure reveals that Training-Time Proxy Integration sits adjacent to Decoding-Time Proxy Tuning (which applies proxies only at inference) and Proxy-Based Architecture Search. Neighboring branches include Knowledge-Guided Adaptation Strategies, which emphasize structured knowledge integration and domain-specific adaptation, and Parameter-Efficient Adaptation Strategies, focusing on low-rank updates and modular adapters. KnowProxy bridges these areas by combining proxy-based efficiency with knowledge-guided steering, diverging from purely mechanistic proxy methods by incorporating explicit reasoning elicitation and from pure knowledge integration approaches by maintaining the proxy architecture paradigm.

Among 30 candidate papers examined, none were identified as clearly refuting any of the three core contributions: the KnowProxy framework itself, the dynamic routing mechanism, and the knowledge elicitation process. Each contribution was assessed against 10 candidates with zero refutable overlaps found. This suggests that within the limited search scope, the combination of knowledge-guided proxy training without probability distribution access appears relatively unexplored. However, the modest search scale (30 candidates from semantic search) means the analysis captures immediate neighbors rather than exhaustive prior work, and the absence of refutations reflects this bounded examination rather than definitive novelty.

Based on the limited literature search, the work appears to occupy a distinctive position combining proxy-based efficiency with knowledge-guided adaptation. The taxonomy context shows this sits at the intersection of two active research threads, and the contribution-level analysis found no direct overlaps among examined candidates. However, the 30-paper search scope and the presence of two sibling papers in the same taxonomy leaf suggest caution in claiming broad novelty without deeper investigation of related proxy tuning and knowledge distillation literature.

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

Contribution 1: KnowProxy framework for knowledge-guided proxy adaptation

Description: The authors introduce a novel proxy-based fine-tuning framework that adapts large language models by training smaller proxy models on textual knowledge and reasoning elicited from frozen LLMs, rather than relying on probability distributions. This design enables applicability to black-box settings where only text outputs are available.

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. LVLM-HBA: Large Vision-Language Model with Cross-Modal Alignment for Human Behavior Analysis

URL: [View paper](#)

Brief Assessment

LVLM-HBA[63] focuses on vision-language models for human behavior analysis using LoRA for parameter-efficient tuning, not on proxy-based fine-tuning of LLMs using textual knowledge.

2. Llms as repositories of factual knowledge: Limitations and solutions

URL: [View paper](#)

Brief Assessment

Factual Knowledge Limitations[60] focuses on evaluating LLMs as knowledge repositories and proposes entity-aware fine-tuning (ENAF) for structured entity representations. It does not address proxy-based fine-tuning frameworks or training proxy models on textual knowledge from frozen LLMs.

3. Large language model-assisted surrogate modelling for engineering optimization

URL: [View paper](#)

Brief Assessment

Surrogate Modelling Optimization[62] focuses on using LLMs to assist in selecting and training surrogate models for engineering optimization tasks, not on proxy-based fine-tuning of LLMs using textual knowledge. The technical domains and objectives are fundamentally different.

4. Learning to rewrite: Generalized llm-generated text detection

URL: [View paper](#)

Brief Assessment

Learning to Rewrite[59] focuses on detecting LLM-generated text through rewriting-based edit distance analysis, not on proxy-based fine-tuning of LLMs using textual knowledge for task adaptation.

5. Tuning Language Models by Proxy

URL: [View paper](#)

Brief Assessment

Tuning by Proxy[6] focuses on reweighting probability distributions from black-box LLMs using smaller proxy models, while the original paper uses textual knowledge and reasoning instead of distributions. These are fundamentally different proxy-based approaches with distinct technical mechanisms.

6. Small models, big insights: Leveraging slim proxy models to decide when and what to retrieve for llms

URL: [View paper](#)

Brief Assessment

Slim Proxy Models[57] focuses on using a proxy model to determine when and what to retrieve for RAG systems, not on proxy-based fine-tuning using textual knowledge from LLMs. The candidate addresses retrieval necessity judgment and query formulation, while the original contribution concerns adapting LLMs through proxy models trained on elicited textual knowledge and reasoning.

7. FedPT: federated proxy-tuning of large language models on resource-constrained edge devices

URL: [View paper](#)

Brief Assessment

FedPT[58] focuses on federated learning scenarios where devices collaboratively tune smaller models for privacy-preserving fine-tuning of black-box LLMs, whereas the original paper addresses knowledge-guided proxy adaptation using textual reasoning from frozen LLMs in non-federated settings.

8. Latuner: An llm-enhanced database tuning system based on adaptive surrogate model

URL: [View paper](#)

Brief Assessment

Latuner[56] focuses on database knob tuning using LLMs to guide Bayesian optimization for configuration selection, not on proxy-based fine-tuning of language models through textual knowledge adaptation.

9. Constructing surrogates for atomistic simulations via deep learning and generative large language models

URL: [View paper](#)

Brief Assessment

Surrogates Atomistic Simulations[55] focuses on constructing surrogate models for atomistic simulations using deep learning (DSO, DNN) and generative LLMs (GPT-4o) to predict material properties. It does not address proxy-based fine-tuning of language models using textual knowledge, which is the core contribution of the original paper.

10. LLM-Based Adaptive Distribution Voltage Regulation Under Frequent Topology Changes: An In-Context MPC Framework

URL: [View paper](#)

Brief Assessment

Adaptive Distribution Voltage[61] addresses distribution voltage regulation in power systems using LLM-based in-context learning for topology adaptation, not proxy-based fine-tuning of language models with textual knowledge.

Contribution 2: Dynamic routing mechanism for adaptive proxy invocation

Description: The authors develop an adaptive routing mechanism that uses uncertainty scores elicited from the LLM's generated knowledge to determine when to invoke the proxy model. This allows the framework to selectively engage the proxy only for uncertain or unreliable LLM outputs, reducing inference overhead while maintaining accuracy.

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. Cargo: A framework for confidence-aware routing of large language models

URL: [View paper](#)

Brief Assessment

Cargo[49] focuses on routing between multiple complete LLMs based on confidence scores to select which model answers a query. The original paper's contribution involves routing to determine when to invoke a proxy model that adapts a frozen LLM's outputs, which is a fundamentally different architectural approach.

2. At-cxr: Uncertainty-aware agentic triage for chest x-rays

URL: [View paper](#)

Brief Assessment

At-cxr[52] focuses on medical imaging triage with uncertainty-based routing for chest X-rays, not general LLM adaptation through proxy models. The routing mechanism serves a different purpose (clinical triage decisions) in a different domain (medical imaging) compared to the original paper's LLM fine-tuning framework.

3. Confident or Seek Stronger: Exploring Uncertainty-Based On-device LLM Routing From Benchmarking to Generalization

URL: [View paper](#)

Brief Assessment

Uncertainty Based Routing[45] focuses on routing between small language models (SLMs) and LLMs based on uncertainty scores for on-device deployment efficiency. The original paper's routing mechanism adaptively invokes a proxy model trained on LLM-generated textual knowledge, which is a fundamentally different architecture and purpose.

4. Program Arrives Home Smoothly: Uncertainty-Based Routing Scheduling of Home-Based Elderly Care Programs

URL: [View paper](#)

Brief Assessment

Home Care Routing[54] addresses routing scheduling for elderly care services using robust optimization under uncertainty, not adaptive model invocation based on LLM uncertainty scores. The routing mechanism here pertains to physical service delivery routes, not computational proxy model activation.

5. Inclusive prompt engineering for large language models: a modular framework for ethical, structured, and adaptive AI

URL: [View paper](#)

Brief Assessment

Inclusive Prompt Engineering[50] focuses on prompt-level strategies for LLMs without model weight modification, using uncertainty-driven prompt selection and persona-based filtering. The original paper's dynamic routing mechanism operates at the model invocation level, selectively engaging a trained proxy model based on LLM output uncertainty scores—a fundamentally different architectural approach not addressed in the candidate.

6. DynMoLE: Boosting Mixture of LoRA Experts Fine-Tuning with a Hybrid Routing Mechanism

URL: [View paper](#)

Brief Assessment

DynMoLE[47] focuses on routing mechanisms for Mixture of LoRA Experts in parameter-efficient fine-tuning, using entropy-based dynamic routing to select experts. The original paper's contribution addresses adaptive proxy model invocation based on LLM uncertainty scores for knowledge-guided adaptation, which is a fundamentally different architectural approach and application context.

7. Closing the Data Loop: Real-World AUVs Adaptive Sampling for Improved Ocean Model Predictions

URL: [View paper](#)

Brief Assessment

Adaptive Sampling AUVs[53] addresses autonomous underwater vehicle path planning for ocean data collection, not LLM adaptation or proxy model invocation based on uncertainty scores.

8. Learning to route llms with confidence tokens

URL: [View paper](#)

Brief Assessment

Confidence Tokens Routing[51] focuses on routing queries between different-sized LLMs based on confidence tokens, not on adaptive proxy model invocation for knowledge-guided fine-tuning as in the original paper's framework.

9. Mediator: Memory-efficient llm merging with less parameter conflicts and uncertainty based routing

URL: [View paper](#)

Brief Assessment

Mediator[48] focuses on model merging with uncertainty-based expert routing for combining multiple fine-tuned LLMs, not on adaptive proxy invocation during LLM fine-tuning. The routing mechanism in Mediator[48] selects between merged experts for inference, whereas the original paper's contribution involves determining when to invoke a proxy model during training/adaptation based on LLM output uncertainty.

10. Collm: Industrial large-small model collaboration with fuzzy decision-making agent and self-reflection

URL: [View paper](#)

Brief Assessment

Collm[46] focuses on industrial large-small model collaboration with fuzzy decision-making for sample routing, while the original paper develops uncertainty-based routing specifically for proxy model invocation in LLM adaptation. The technical approaches and application contexts differ substantially.

Contribution 3: Knowledge elicitation and filtering process for proxy training

Description: The authors propose a method to extract textual knowledge and reasoning from LLMs via prompting, along with confidence scores for each piece of knowledge. A filtering process retains only high-confidence knowledge, which is then used to train the proxy model to align LLM-derived reasoning with target task distributions.

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. Knowledge Elicitation with Large Language Models for Interpretable Cancer Stage Identification from Pathology Reports

URL: [View paper](#)

Brief Assessment

Knowledge Elicitation Cancer[40] focuses on extracting staging rules from pathology reports for cancer classification, not on training proxy models to adapt LLMs. The knowledge elicitation serves a different purpose (direct inference for medical staging) rather than proxy-based LLM adaptation.

2. KnowledgePrompts: Exploring the Abilities of Large Language Models to Solve Proportional Analogies via Knowledge-Enhanced Prompting

URL: [View paper](#)

Brief Assessment

KnowledgePrompts[36] focuses on proportional analogy completion tasks using knowledge-enhanced prompts for direct LLM inference, not on proxy model training or knowledge filtering for adaptation purposes.

3. MoRE-LLM: Mixture of Rule Experts Guided by a Large Language Model

URL: [View paper](#)

Brief Assessment

MoRE-LLM[38] focuses on rule extraction and refinement from LLMs for interpretable classification, not on knowledge elicitation with confidence scores for proxy model training in reasoning tasks.

4. Efficient Knowledge Probing of Large Language Models by Adapting Pre-trained Embeddings

URL: [View paper](#)

Brief Assessment

Efficient Knowledge Probing[42] focuses on using pre-trained embeddings to predict LLM knowledge without querying the model, rather than eliciting textual knowledge through prompting for proxy training. The candidate's approach adapts embedding models to match LLM outputs, while the original paper extracts knowledge via prompting and filters it by confidence scores for training proxies.

5. Accelerated Preference Elicitation with LLM-Based Proxies

URL: [View paper](#)

Brief Assessment

Accelerated Preference Elicitation[35] focuses on eliciting preferences from bidders in combinatorial auctions using natural language, not on extracting textual knowledge from LLMs for proxy model training in reasoning tasks.

6. Accelerating Catalysis Understanding via Large Language Model Data Extraction and Shallow Machine Learning Techniques

URL: [View paper](#)

Brief Assessment

Catalysis Data Extraction[41] focuses on extracting textual data from scientific literature for catalysis research using LLMs, not on eliciting knowledge through prompting for proxy model training or fine-tuning language models.

7. Self-AMPLIFY: Improving Small Language Models with Self Post Hoc Explanations

URL: [View paper](#)

Brief Assessment

Self-AMPLIFY[44] focuses on generating rationales from post hoc explanations applied to the same small model to improve its own performance via in-context learning, rather than eliciting knowledge from frozen LLMs to train separate proxy models as in the original paper.

8. A Tale of LLMs and Induced Small Proxies: Scalable Agents for Knowledge Mining

URL: [View paper](#)

Brief Assessment

Induced Small Proxies[43] focuses on knowledge mining tasks where LLMs annotate data to train small proxies for entity extraction and classification. The original paper's contribution centers on eliciting textual knowledge with confidence scores for proxy training in reasoning tasks, which is a different application domain and methodology.

9. Large Language Model as Meta-Surrogate for Data-Driven Many-Task Optimization: A Proof-of-Principle Study

URL: [View paper](#)

Brief Assessment

Meta Surrogate Optimization[37] focuses on using LLMs as meta-surrogates for fitness prediction in many-task optimization, not on proxy model training with knowledge elicitation and filtering for LLM adaptation as in the original paper.

10. A Self-feedback Knowledge Elicitation Approach for Chemical Reaction Predictions

URL: [View paper](#)

Brief Assessment

Self Feedback Elicitation[39] focuses on chemical reaction predictions using self-feedback clustering for reaction type annotation, not on proxy model training for LLM adaptation through knowledge elicitation and filtering.

Appendix: Text Similarity Detection

Textual similarity detection checked 32 papers and found 2 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. Tuning Language Models by Proxy

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] KnowProxy: Adapting Large Language Models by Knowledge-guided Proxy [View paper](#)
- [1] KaSA: Knowledge-Aware Singular-Value Adaptation of Large Language Models [View paper](#)
- [2] Bridging Stepwise Lab-Informed Pretraining and Knowledge-Guided Learning for Diagnostic Reasoning [View paper](#)
- [3] KG-SR-LLM: Knowledge-Guided Semantic Representation and Large Language Model Framework for Cross-Domain Bearing Fault Diagnosis [View paper](#)
- [4] KN-VLM: KNowledge-guided Vision-and-Language Model for visual abductive reasoning [View paper](#)
- [5] Knowledge and Task-Driven Multimodal Adaptive Transfer Through LLMs with Limited Data [View paper](#)
- [6] Tuning Language Models by Proxy [View paper](#)
- [7] Adapting Large Language Models to Log Analysis with Interpretable Domain Knowledge [View paper](#)
- [8] Enhancing Multimodal LLM for Detailed and Accurate Video Captioning using Multi-Round Preference Optimization [View paper](#)

- [9] USTC-BUPT at SemEval-2024 Task 8: Enhancing Machine-Generated Text Detection via Domain Adversarial Neural Networks and LLM Embeddings [View paper](#)
- [10] Hierarchical cross-modal prompt learning for vision-language models [View paper](#)
- [11] Lpzero: Language model zero-cost proxy search from zero [View paper](#)
- [12] Aligning Teacher with Student Preferences for Tailored Training Data Generation [View paper](#)
- [13] Direct Preference Knowledge Distillation for Large Language Models [View paper](#)
- [14] Improving instruction following in language models through proxy-based uncertainty estimation [View paper](#)
- [15] Revolutionizing Mixed Precision Quantization: Towards Training-free Automatic Proxy Discovery via Large Language Models [View paper](#)
- [16] InfuserKI: Enhancing Large Language Models with Knowledge Graphs via Infuser-Guided Knowledge Integration [View paper](#)
- [17] Flipping Knowledge Distillation: Leveraging Small Models' Expertise to Enhance LLMs in Text Matching [View paper](#)
- [18] Continuously Steering LLMs Sensitivity to Contextual Knowledge with Proxy Models [View paper](#)
- [19] OpenAVS: Training-Free Open-Vocabulary Audio Visual Segmentation with Foundational Models [View paper](#)
- [20] Learn as Individuals, Evolve as a Team: Multi-agent LLMs Adaptation in Embodied Environments [View paper](#)
- [21] FedPromo: Federated Lightweight Proxy Models at the Edge Bring New Domains to Foundation Models [View paper](#)
- [22] Selective Knowledge Injection via Adapter Modules in Large-Scale Language Models [View paper](#)
- [23] Large and Small Model Collaboration for Air Interface [View paper](#)
- [24] FedLMA: A Federated Learning Framework Integrating LLM-Based Multi-Agent Reasoning With Knowledge Distillation [View paper](#)
- [25] Augmented Large Language Models with Parametric Knowledge Guiding [View paper](#)
- [26] KgPLM: Knowledge-guided Language Model Pre-training via Generative and Discriminative Learning [View paper](#)
- [27] Large language model uncertainty proxies: discrimination and calibration for medical diagnosis and treatment. [View paper](#)
- [28] Towards Mitigating Excessive Forgetting in LLM Unlearning via Entanglement-Aware Unlearning with Proxy Constraint [View paper](#)
- [29] Deciphering Invariant Feature Decoupling in Source-free Time Series Forecasting with Proxy Denoising [View paper](#)
- [30] NoEsis: Differentially Private Knowledge Transfer in Modular LLM Adaptation [View paper](#)
- [31] A Multi-Modal Knowledge-Enhanced Framework for Vessel Trajectory Prediction [View paper](#)
- [32] Incorporating Patient Values in Large Language Model Recommendations for Surrogate and Proxy Decisions. [View paper](#)
- [33] Proxy Denoising for Source-Free Domain Adaptation [View paper](#)
- [34] Self-Knowledge Guided Retrieval Augmentation for Large Language Models [View paper](#)
- [35] Accelerated Preference Elicitation with LLM-Based Proxies [View paper](#)
- [36] KnowledgePrompts: Exploring the Abilities of Large Language Models to Solve Proportional Analogies via Knowledge-Enhanced Prompting [View paper](#)
- [37] Large Language Model as Meta-Surrogate for Data-Driven Many-Task Optimization: A Proof-of-Principle Study [View paper](#)
- [38] MoRE-LLM: Mixture of Rule Experts Guided by a Large Language Model [View paper](#)
- [39] A Self-feedback Knowledge Elicitation Approach for Chemical Reaction Predictions [View paper](#)
- [40] Knowledge Elicitation with Large Language Models for Interpretable Cancer Stage Identification from Pathology Reports [View paper](#)
- [41] Accelerating Catalysis Understanding via Large Language Model Data Extraction and Shallow Machine Learning Techniques [View paper](#)
- [42] Efficient Knowledge Probing of Large Language Models by Adapting Pre-trained Embeddings [View paper](#)
- [43] A Tale of LLMs and Induced Small Proxies: Scalable Agents for Knowledge Mining [View paper](#)
- [44] Self-AMPLIFY: Improving Small Language Models with Self Post Hoc Explanations [View paper](#)
- [45] Confident or Seek Stronger: Exploring Uncertainty-Based On-device LLM Routing From Benchmarking to Generalization [View paper](#)
- [46] Collm: Industrial large-small model collaboration with fuzzy decision-making agent and self-reflection [View paper](#)
- [47] DynMoLE: Boosting Mixture of LoRA Experts Fine-Tuning with a Hybrid Routing Mechanism [View paper](#)
- [48] Mediator: Memory-efficient llm merging with less parameter conflicts and uncertainty based routing [View paper](#)
- [49] Cargo: A framework for confidence-aware routing of large language models [View paper](#)
- [50] Inclusive prompt engineering for large language models: a modular framework for ethical, structured, and adaptive AI [View paper](#)
- [51] Learning to route llms with confidence tokens [View paper](#)
- [52] At-cxr: Uncertainty-aware agentic triage for chest x-rays [View paper](#)
- [53] Closing the Data Loop: Real-World AUVs Adaptive Sampling for Improved Ocean Model Predictions [View paper](#)
- [54] Program Arrives Home Smoothly: Uncertainty-Based Routing Scheduling of Home-Based Elderly Care Programs [View paper](#)
- [55] Constructing surrogates for atomistic simulations via deep learning and generative large language models [View paper](#)
- [56] Latuner: An llm-enhanced database tuning system based on adaptive surrogate model [View paper](#)
- [57] Small models, big insights: Leveraging slim proxy models to decide when and what to retrieve for llms [View paper](#)
- [58] FedPT: federated proxy-tuning of large language models on resource-constrained edge devices [View paper](#)
- [59] Learning to rewrite: Generalized llm-generated text detection [View paper](#)
- [60] LLMs as repositories of factual knowledge: Limitations and solutions [View paper](#)
- [61] LLM-Based Adaptive Distribution Voltage Regulation Under Frequent Topology Changes: An In-Context MPC Framework [View paper](#)
- [62] Large language model-assisted surrogate modelling for engineering optimization [View paper](#)
- [63] LVLM-HBA: Large Vision-Language Model with Cross-Modal Alignment for Human Behavior Analysis [View paper](#)