

SOLAR: A Self-Optimizing Open-Ended Autonomous Agent for Lifelong Learning and Continual Adaptation

Nitin Vetcha^{1,2,*}, Dianbo Liu¹

¹Department of Ophthalmology, Yong Loo Lin School of Medicine, National University of Singapore, Singapore

²Department of Computational and Data Sciences, Indian Institute of Science, Bangalore, Karnataka, India

Abstract

Despite the remarkable success of large language models (LLMs), they still face bottlenecks while deploying in dynamic, real-world settings with primary challenges being concept drift and the high cost of gradient-based adaptation. Traditional fine-tuning (FT) struggles to adapt to non-stationary data streams without resulting in catastrophic forgetting or requiring extensive manual data curation. To address these limitations within the streaming and continual learning paradigm, we propose the Self-Optimizing Lifelong Autonomous Reasoner (SOLAR) which is an open-ended autonomous agent that leverages parameter-level meta-learning to self-improve, treating model weights as an environment for exploration. It initiates the process by consolidating a strong prior over common-sense knowledge making it effective for transfer-learning. By utilizing a multi-level reinforcement learning approach, SOLAR autonomously discovers adaptation strategies, enabling efficient test-time adaptation to unseen domains. Crucially, SOLAR maintains an evolving knowledge base of valid modification strategies, implicitly acting as an episodic memory buffer to balance plasticity (adaptation to new tasks) and stability (retention of meta-knowledge). Experiments demonstrate that SOLAR outperforms strong baselines on common-sense, mathematical, medical, coding, social and logical reasoning tasks, marking a significant step toward autonomous agents capable of lifelong adaptation in evolving environments.

Keywords

Continual Adaptation, Lifelong Learning, Self-Evolution, Test-Time Adaptation, Transfer-Learning, Large Language Models

1. Introduction

Large Language Models (LLMs) possess remarkable emergent abilities due to massive pretraining. However, deploying them in streaming environments reveals a critical weakness which is the inability to adapt to non-stationary data distributions (concept drift) without expensive retraining or human intervention. While Parameter-Efficient Fine-Tuning (PEFT) techniques like LoRA (Hu et al. 2022) reduce the parameter update volume, they remain static solutions that do not inherently address the stability-plasticity dilemma central to Continual Learning (CL). Existing adaptation strategies often rely on generic, hand-crafted heuristics that fail to generalize across the shifting temporal dependencies of real-world streams. This disconnect necessitates a system that can not only adapt parameters on the fly but also learn how to adapt based on accumulating experience. We propose that the high-dimensional weight space of an LLM contains rich meta-knowledge that, if navigated autonomously, can yield bespoke adaptation strategies for novel tasks. This motivates our primary research question:

RQ: Can LLMs learn to modify their internal representation space autonomously to handle concept drift, analogous to how humans assimilate and restructure knowledge in lifelong learning scenarios?

To answer this, we investigate the cognitive science of lifelong learning. As humans, we do not merely memorize new data, instead we restructure our internal schematics in order to be able to

1st Streaming Continual Learning Bridge at AAAI26, January 21, 2026, Singapore.

*Corresponding author.

✉ nitinvetcha@iisc.ac.in (N. Vetcha); dianbo@nus.edu.sg (D. Liu)

🌐 <https://github.com/nitinvetcha/> (N. Vetcha); <https://www.asintelligence.xyz/> (D. Liu)

🆔 0009-0003-6542-324X (N. Vetcha); 0000-0002-3042-9161 (D. Liu)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

accommodate new information while simultaneously retaining prior heuristics. This process is what has essentially enabled humans to navigate non-stationary environments. For instance, a student adapts their study strategy based on the nature of a new subject (plasticity) without unlearning how to study generally (stability). Current LLM adaptation, by contrast, is often rigid since models consume task data “as-is”, failing to develop bespoke internal transformation strategies. To replicate this cognitive flexibility, we introduce SOLAR (Self-Optimizing Lifelong Autonomous Reasoner). It functions as a meta-learning agent that decouples rapid task adaptation (streaming machine learning) from long-term strategy retention (continual learning). By discovering and validating parameter-level modifications, SOLAR enables efficient adaptation to unseen tasks while populating a persistent knowledge base to mitigate catastrophic forgetting. This work thus bridges the gap between static parameter generation and dynamic, lifelong self-evolution. Further more, by grounding the search space in neural network weights, we target generalized principles of model capability rather than task-specific memorization. Just as scaling laws (Kaplan et al. 2020) predict performance based on size, we posit that predictable weight-modification patterns exist that allow for rapid, data-efficient adaptation to concept drift, minimizing the lag between detecting a distributional shift and deploying an updated model. The remainder of this paper is organized as follows. In Section 2, we highlight the motivation for our approach in detail. Section 3 presents the literature survey conducted, Section 4 contains the methodology with implementation specifics in Section 5. Experimental results are provided in Section 6 and in Section 7, we present our concluding remarks.

2. Motivation

Our primary motivation stems from human psychology and pedagogy. For example, consider a human student who is preparing to take an end-sem examination of a machine learning course. Quite often, students tend to rely on their prior prepared notes for preparation. These notes are often derived from the lecture content, textbooks or information available on the internet. Thus instead of relying on the raw content, students assimilate and rewrite the information in the form of notes as per their own intrinsic reasoning skill and aptitude. This improves the capability of students to comprehend the content better and therefore respond well to the exam questions. This phenomenon of reinterpreting and augmenting external knowledge in a way that is easier to understand as well as developing the necessary skill-sets is not limited to just taking exams, but seems to be universally true of human learning across tasks. Furthermore, depending on one’s interests, humans assimilate information in different ways - some might condense the information into a visual diagram, some into text, or some might rely more on concrete mathematical descriptions. Such restructuring or development of internal knowledge as well as assimilation or rewriting of external information, as part of the learning process is in contrast with how LLMs undergo currently training and adaptation. Given a new task, current LLMs consume and learn from the task data “as-is” via finetuning or in-context learning. The issue with this, just like in the human setting, such data may not be in an optimal format (or volume) for learning, or there might not be the relevant skill-set developed to learn it and current approaches do not enable models to develop bespoke strategies for how to best transform themselves internally or even learn from their training data. In this work, we therefore investigate the question as to if it is possible for even LLMs, analogous to humans, to suggest strategies by themselves which can enable them to perform better on a given task.

A secondary source of motivation as to why we ground our strategy search space in the neural network weights is because unlike task-specific knowledge, the weight-level meta-knowledge represents generalized principles about how neural network parameters relate to model capabilities, thereby providing crucial insights for self-evolving agents. There are several prior research works which have shown that there exists a positive correlation between types of neural network weight patterns and downstream model performance characteristics. For example, scaling laws research [1] has demonstrated that there are predictable relationships between model size and performance. Similarly, structured

sparsity learning gives an indication so as to how particular weight patterns can be useful for developing more efficient representations [2].

3. Related Work

Test-Time Training (TTT) is a recently emerging class of approaches which updates model weights at inference time using techniques such as input perplexity or cross-entropy minimization on only unlabeled test data enabling self supervised enhancement of LLM performance [3, 4] or via reinforcement learning by utilizing the priors in the pre-trained models [5] or by using reflection and verifier-driven sample selection [6, 7] or by using a task-specific curriculum [8] or by using a mixture-of-expert based model merging [9]. An alternative approach is to scale inference compute at test time as well using for example ensemble approaches such as majority voting. While test-time approaches is a promising option, such a computational overhead might not be necessary always and it often fails in cases where data is scarce or quality of unlabeled data is poor.

Adversarial Fine Tuning is another emerging class of techniques where in two LLM instances are made to debate with each other about a topic or one instance serves as a challenger or teacher and the other instance serves as a solver or student to generate synthetic data, either from unlabeled prompts or even from scratch itself and use approaches like majority voting to create pseudo-labels which can further be used for updating model's knowledge accordingly [10, 11, 12]. This can also be done by some additional fine-tuning using information which is available in the LLM's context as well [13] similar to knowledge distillation. Recent works include SQLM [14], R-Zero [15], TT-SI [16], SIRLC [17]. While this is an efficient approach in data scarce domains where TTT fails, it is not always efficient as there are certain challenging domains which require mastering novel reasoning skills and it is well known that scaling data isn't sufficient in this regimes such as mathematics [18].

Reinforcement Learning (RL) is a well established approach for pushing the capabilities of LLMs and recent works such as SEAL [19], RLAIF [20], SRLM [21] and Memento [22], which uses a memory-based online RL policy have shown promising potential in the low-cost continual adaptation of LLMs. In RL, **meta-learning** has been used as well in order train agents in scenarios where it needs to learn novel tasks quickly [23]. SOLAR can be seen as thus following meta-learning principles since it learns an adaptation strategy i.e., how to generate effective self weight update using a meta optimization loop. Closely, related are **self-referential** systems as well which learn to update their own parameters as in [24] and **self-evolving** agents which enable LLM to improvise by autonomously acquiring, refining and learning from experiences generated by the model itself [25, 26]. While RL based approaches are quite good, its often challenging to achieve convergence and design optimal policies which are efficient in terms of compute and time as well.

Parameter Generation is another research direction which has seen several pioneering works such as RPG [27], DnD [28], T2L [29], ORAL [30], COND P-DIFF [31]. DnD generates task-specific parameters from unlabeled prompts without per-task training via a prompt-conditioned hyper-convolutional decoder while T2L does the same but uses a hyper-network and task description instead. ORAL leverages architectural and textual conditioning for flexible, scalable LoRA parameter adaptation. RPG introduces a recurrent diffusion architecture for scalable unconditional LoRA parameter generation. COND P-DIFF applies conditional latent diffusion for controllable LoRA parameter synthesis with strong cross-domain generalization. An associated direction is *model merging* as well, which facilitates generalization to unseen tasks via **multi-task learning** [32, 33]. While these works have been effective, the limitation is that these are static parameters which once generated do not undergo any further modification but this feature is crucial for domains requiring the implicit meta-knowledge.

4. Methodology

In this section, we describe the framework of our proposed approach (see Figure 1). SOLAR starts by treating the LLMs own weights as environment variables to explore, upon which it would systematically propose scientific hypotheses to modify the internal representation space appropriately so as to adapt the LLM to the unseen task. A major challenge for the design, therefore, is the high dimensionality and non-convexity of the LLM weight space itself which makes the initialization and subsequent exploration process extremely complex. To overcome this, we work only with low-rank parameters [34] which constitutes a much smaller fraction ($\sim 1\%$) of the original model’s weights. In addition, to avoid the limitations arising from selecting a single starting point, which might not be optimal to wiggle around, we prefer to sample from a plausible weight distribution space. This step is essential to eliminate the risk of non-convergence. To get this initial distribution for weights i.e., self-weight sampling, we refer to prior works in large-scale LLM parameter generation and use a convolution-based decoder architecture as the backbone for SOLAR’s exploration point initializer.

Once the weights have been initialized¹ for exploration, SOLAR then uses a foundation-model-based agent, which is for now simply an LLM trained using reinforcement learning (RL) to come up with probable hypotheses at inference time for weight-space exploration using test-time scaling and compute. To however, facilitate the training process, it is necessary to first curate by hand a seed knowledge base, consisting of either proven or plausible weight modification strategies, which will then serve as the action space for LLM’s initial stages of exploration during RL training. This would be a multi-stage recipe consisting of three distinct progressively harder levels. Level I consists of training the LLM to produce only single valid and efficient self-edits (a self-edit as the name suggests is basically a modification strategy proposed by an LLM to update its own weights depending on the task) from among the ones present in the initial knowledge base. Level II comprises of training the model to output chain-of-self-edits, since coupling strategies sequentially is also helpful (moreover if viewed in a abstract sense, it can be considered in effect as a single complex edit which can be decomposed into simpler instances). Level III is a significantly challenging aspect both for the LLMs as well as from implementation perspective as well, which is basically letting LLMs to explore the hypothesis space in its entirety, thereby going beyond human-crafted approaches. A positive performance in Level III would be a significant leap as it could possibly open up new frontiers in training and fine-tuning paradigms as has been similarly done in other areas as well such as neural architecture search [36] and optimization [37].

After plausible hypothesis have been generated by the foundation model-based agent and implemented, its necessary to test the hypothesis. For this purpose, we create a separate evaluation split if available. However, since SOLAR is designed to adapt LLMs efficiently to unseen tasks as well, the dataset for evaluation itself would be generated on the fly using adversarial approaches involving multiple instances of an LLM, one proposing and one solving questions on a particular topic as in SQLM [14] or R-Zero [38]. Once the hypothesis has been tested and is found to be valid (as in it improves performance in some pre-determined metric such as accuracy on the eval set), it would be then added back into the knowledge base, thereby enriching the action space of LLM for future iterations. In order to prevent catastrophic forgetting, SOLAR implements a meta-level weight regularization technique as well. Therefore, by automating the process of self-improvement using principled methodologies and meta-knowledge in a scientific manner (i.e., propose, validate and accept hypotheses), SOLAR provides a holistic framework towards the next generation of **AI generating AI** agents, because as soon as web-scale data corpora is exhausted, progress will hinge on a model’s capacity to generate its own high-utility training signal.

¹These weights can optionally be encoded into a structured representation correlated with network performance like world models such as JEPA [35].

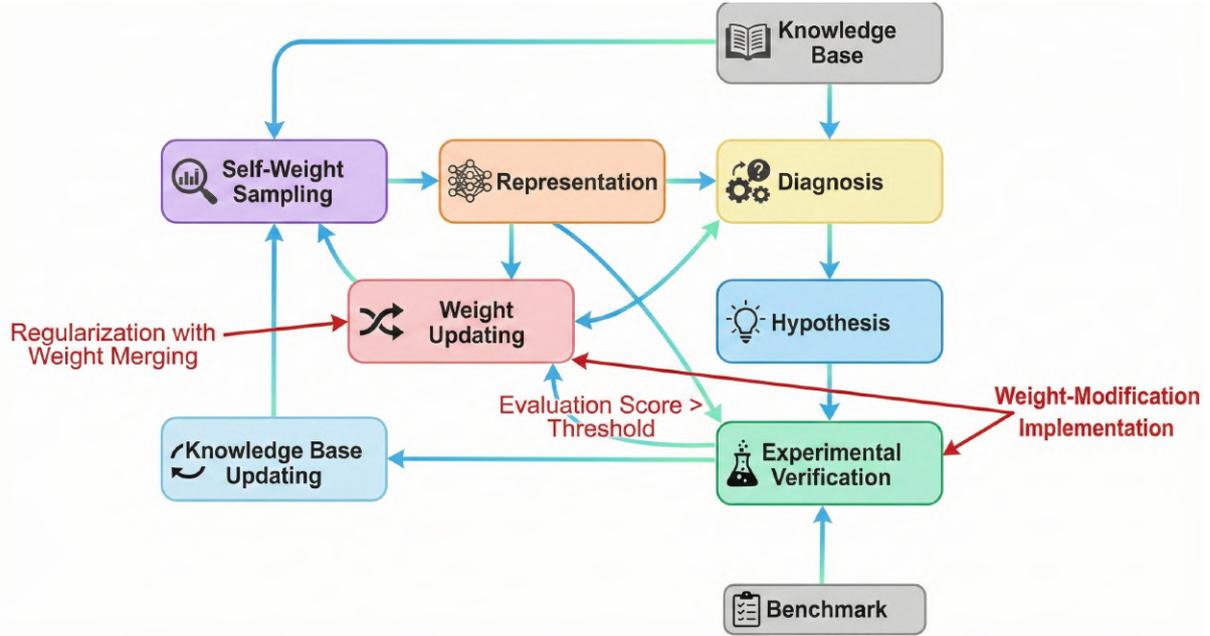


Figure 1: SOLAR’s methodology of weight-level meta-knowledge discovery and modification summarized

5. Implementation

5.1. Architecture

Primary architectural detail in SOLAR’s framework is the design of the weight-space exploration initializer. As mentioned in Section 4, we use a convolution based decoder model for this purpose. We assume that we have access to either the unseen task’s description or atleast a handful of unlabeled examples representative of its requirements. We then send them to an open-sourced text encoder for embedding extraction. This extraction process can be formally represented as, $c_i = \text{Encoder}(p_i, \theta)$, where $\text{Encoder}(\cdot, \cdot)$ denotes the embedding extraction function parameterized by θ , and c_i represents the extracted embedding corresponding to prompt p_i . We use an encoder-based language model architecture for this purpose i.e., Sentence-BERT (all-MiniLM-L6-v2 specifically) [39]².

Next, following [27], is the parameter tokenization process (see Figure 2), which is done so as to preserve both the layer-wise distribution and the cross-layer correlations. Specifically, (i) weights are split according to their layer indices, (ii) layer-wise normalization is applied to mitigate distribution shifts, (iii) parameters are sliced into non-overlapping tokens with uniform size, and (iv) a lightweight permutation state (encoded as a one-hot vector) is used to alleviate symmetry issues [40] when collecting multiple checkpoints. Additionally, 2D position embeddings (first dimension encodes layer index, while second dimension captures the token’s in-layer position). [41] are employed to ensure the network retains positional awareness of each token within the entire set. In our case, each LoRA matrix is of shape 8×896 , which is then split into 7 smaller chunks, each with a shape of 8×128 , which is then finally padded to a uniform size of 10×130 .

Say, the dimension of prompt embeddings is $[B, N, L, C]$ where B, N, L and C denote batch size, length of prompt batch (i.e., number of prompts), sequence length, and hidden dimension, respectively. The decoder (see Figure 3) consists of multiple sequential layers, each performing 5 2D convolutions. These convolutions are divided into three categories: i) **width convolution** that operates on (C, L)

²It is to be noted that BERT’s supported sequence length is only 512 and for longer sequences, padding should be done. However, in our use case, maximum sequence length is only 384 and thus padding is not necessary.

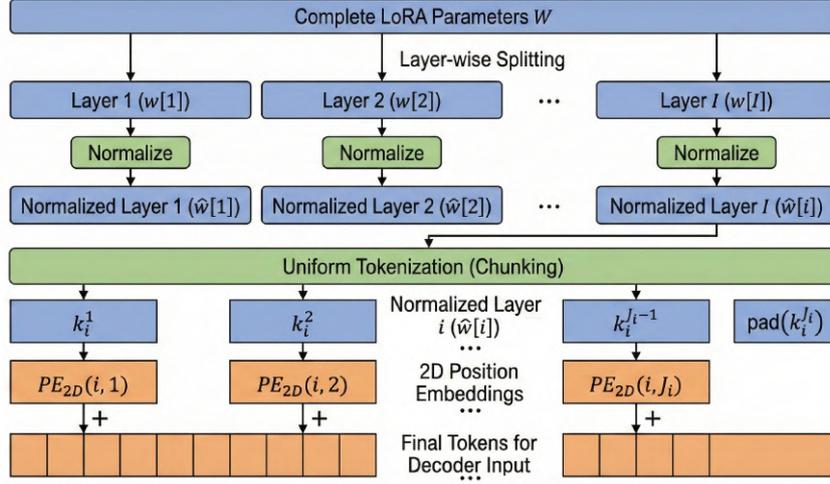


Figure 2: Details of the Parameter Tokenization Process

dimension, ii) **height convolution** that operates on (L, N) dimension) iii) **layer-wise convolution** that on (N, L) dimension), with notations Conv_W , Conv_H , and Conv_L . Each layer consists of two Conv_W , two Conv_H and one Conv_L . Given this, the forward operation of the decoder block is,

$$\begin{aligned}
 c_W^l &= \text{Conv}_H^1(\text{Conv}_W^1(c^{l-1})) \\
 c_H^l &= \text{Conv}_W^2(\text{Conv}_H^2(c^{l-1})) \\
 c^l &= \text{Conv}_L((c_W^l + c_H^l + b)/3)
 \end{aligned}$$

where c^l is hidden state output by the l th layer, c^0 is prompt embedding encoded by the condition extractor, and b is learnable bias. Through this process, input is transformed from dimension $[B, N, L, C]$ to $[B, N', L', C']$ which is then compatible to be converted into a flattened LoRA adapter for the LLM³. In this work, the base LLM used is Qwen2.5-0.5B-Instruct [42] and LoRA is applied to the linear projection layers within both the self-attention mechanism and the MLP blocks of the transformer architecture. Specifically, this includes the query, key, value and output projections in attention blocks, as well as the gate, up and down projections in MLP blocks.

5.2. Training

In this work, we focus on the domain of common-sense reasoning and select 4 datasets for evaluation, namely HellaSwag [43], BoolQ [44] as well as the challenge and easy set of AI2 Reasoning Challenge (ARC) [45]. ARC dataset contains grade-school level, multiple-choice science questions. HellaSwag instructs models to select from choices that best finish the sentence among ground truth and an adversarial set of machine-generated wrong answers. BoolQ is a question answering dataset for yes/no questions containing various factual problems. We use existing checkpoints of these datasets⁴ (batch size was 32 and number of samples was 5000) which have been collected by first pretraining on the target dataset for 75 steps with a learning rate of $1e-4$ and then performing fine-tuning on the target dataset for 50 additional steps with a learning rate of $1e-5$, while saving a checkpoint at each step.

³In our present implementation, the entire flow is $(128,384,384) \rightarrow (128,200,300) \rightarrow (128,100,256) \rightarrow (256,50,200) \rightarrow (512,50,200) \rightarrow (1024,25,200) \rightarrow (1024,10,200) \rightarrow (2048,10,200) \rightarrow (4296,8,128)$

⁴For training however, even Open-Book Question Answering or OBQA [46], Physical Interaction: Question Answering or PIQA [47] and WinoGrande [48] have been used as well. OBQA aims to promote research in advanced question-answering with salient facts summarized as an open book. PIQA focuses on everyday situations with a preference for a typical solutions. WinoGrande features a fill-in-a-blank task with binary options for commonsense reasoning questions.

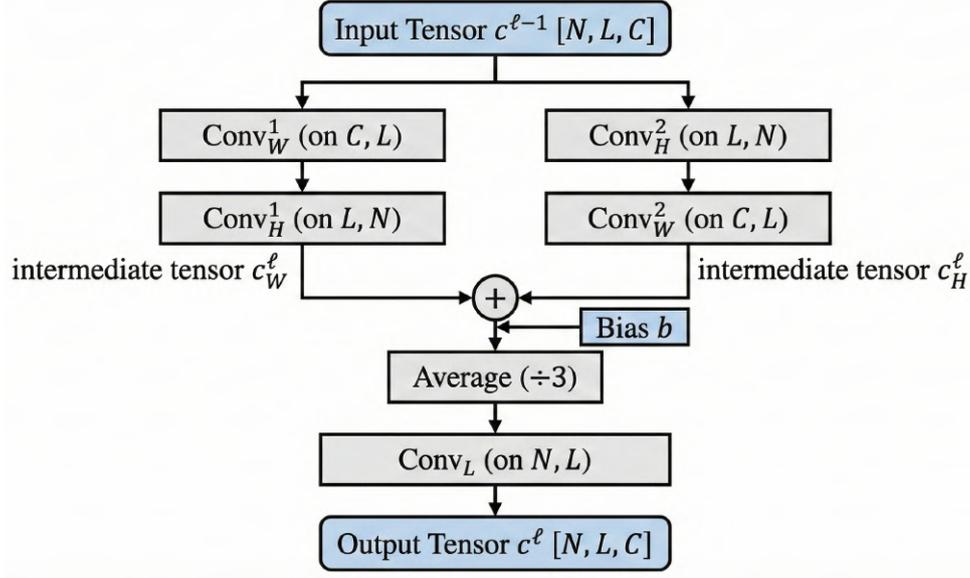


Figure 3: Details of the Hyper-Convolutional Decoder Architecture used

Subsequently, prompt-checkpoint pairing is done as follows. Given a dataset P , it is first divided into non-overlapping prompt batches $[p_1, \dots, p_i, \dots, p_I]$. Denote the trained LLM checkpoints of this dataset as $M = [m_1, \dots, m_j, \dots, m_J]$. Then randomly a batch of prompts and a corresponding checkpoint is picked to create a pair $\{p_i, m_j\}$, which then serves as an input-output data point for training the decoder. The objective function for training is the mean squared error (MSE) loss between the output from the decoder’s last block for a particular prompt batch and the training checkpoint associated with it.

Next crucial step is the hand-crafting of seed knowledge base. To this end, we identify five primary families of strategies⁵, each containing its own sub-strategies as well, namely

- Test-Time Training (TTT) using input perplexity minimization [3] or via reinforcement learning [5] for example by using self-reflection and verification loops like GEPA [49], ReflectEvo [50], REVISE [7] or Instruct-of-Reflection [51]. It could also involve prompt optimization using frameworks like TextGrad [52] or CAST [53]
- Post-training data-free LoRA modifications such as mixing LoRA subspaces obtained by weight decomposition of constituent matrices [54] or bounding norm of selected parameters [55] or evening merging multiple task-specific LoRA adapters [56]
- SQLM [57], R-Zero [15] or SEAL [19] like reinforcement learning based frameworks which enable LLMs to self-adapt by generating their own finetuning data and update directives (another example is TT-SI [16])
- Test-Time Scaling (TTS) using either a router or an ensemble approach i.e., we generate and perform inference with multiple adapters obtained by using different representative prompt batches and to obtain the final prediction, select either the most confident prediction (max_confidence) or by a majority vote or sum_logprobs (i.e., sum log probabilities across adapters per prediction and pick the one with highest total logprob)
- Latent Space (LS) Approaches which aim at working or modifying the internal layers [4] or hidden activations [58] directly of the LLM. It may also involve decoding algorithms which modify the sampling procedure itself [59, 60]. We consider them as part of latent space family because they

⁵Unfortunately, there are no research works highlighting approaches for optimizing the performance of LoRA’s obtained via the process of parameter generation, thereby posing a major challenge in identification of plausible strategies, which had to be cherry-picked via trial and error.

tamper with internal probability distribution of next-tokens unlike other families which modify the parameters explicitly.

We first formulate the objective for outer-loop RL training which generates adaptation strategies AS, as in [19]. Let θ denote the parameters of the language model LM_θ . In order to adapt to an unseen dataset (task) \mathcal{D} , SOLAR requires as specified in Section 4, C which is a context containing information relevant to the task and τ which is the evaluation strategy and metric used to assess the model’s downstream adaptation. Based on C , SOLAR generates an AS and updates its parameters accordingly $\theta' \leftarrow \text{Update}(\theta, \text{AS})$. We thus have an RL setup i.e., the model takes an *action* (generating AS), receives a *reward* r based on $LM_{\theta'}$ ’s performance on τ and updates its policy to maximize expected reward,

$$\mathcal{L}_{\text{RL}}(\theta_t) := -\mathbb{E}_{(C,\tau)\sim\mathcal{D}} \left[\mathbb{E}_{\text{AS}\sim\text{LM}_{\theta_t}(\cdot|C)} [r(\text{AS}, \tau, \theta_t)] \right]$$

It is to be noted that the reward assigned to a given action depends on the model parameters θ at the time the action is taken (since θ is updated to θ' , which is then evaluated). An implication of this is that the while modeling the RL state, one must therefore include θ in the policy’s parameters as well along with C , even though the policy’s observation is limited to C (because it is extremely infeasible to directly place θ in the LLM’s context window). Therefore, the (state, action, reward) triples which have been collected by using an older model weights, θ_{old} , will not be aligned for the current model θ_{current} . Hence, an on-policy approach should be adapted, by which adaptation strategies are sampled from and, even more importantly, the rewards itself will be calculated using the current model.

In particular, the specific on-policy approach used is ReST^{EM} [61] where samples are first generated⁶ from the current model and are filtered by using binary feedback [$r(\text{AS}, \tau, \theta_t)$ is 1 if on τ , AS improves LM_{θ_t} ’s performance and is 0 otherwise]. The model is then fine-tuned on these samples and this continues in an iterative manner (See Algorithm 1).

A subtle detail, which hasn’t yet been covered is the exact nature of the adaptation strategy itself. This depends on the particular strategy family being used, however the format is consistent across all which is basically a JSON object specifying the particular configurations to be used⁷. It contains a field, `family` which takes values TTT, LoRA and TTS. Currently, the following choices have been experimented

- For TTT, we use [3] and the corresponding JSON object has fields `ttl_steps` (number of training steps in the TTL loop), `learning_rate`, `batch_size` and `shuffle_data` (boolean variable).
- For LoRA modifications, we use two-subspace (TS) mixing version from [54] and the corresponding JSON object has only a single field, namely `lambda` which is a hyperparameter determining the ratio in which the two resulting subspaces must be mixed.
- For TTS, we use either an ensemble or router approach. In the router approach (see Figure 4), we basically sample multiple prompt batches and choose that batch whose average of similarity scores⁸ of individual prompts (M1) or averaged prompt embedding (M2), is closest to that of the question at test time. The corresponding JSON object has fields `num_prompt_batches` (indicating the number of prompt batches to be sampled from the test split of unseen dataset) and `method`

⁶Currently, only a deterministic number of samples are being generated, 15 to be precise. This could however be improvised to be dynamic in future version of the work wherein samples would continue to be generated until a particular confidence threshold, as determined by the model itself is reached instead. The same is true for number of iterations as well which is just 2 for now.

⁷Since the model being used is Qwen2.5-0.5B-Instruct, it was facing difficulty in following instructions given in the prompt for generation of structured outputs even after temperature alteration. In such cases, verification and formatting was done by using Qwen2.5-7B-Instruct instead.

⁸Cosine similarity and Euclidean distance were tested and the latter was found to perform better empirically. Thus, `avg_sim_score` and `avg_prompt_embed`. use euclidean distance by default. Alternatively, measure of similarity can also be included as a new field but hasn’t been explored in the current work.

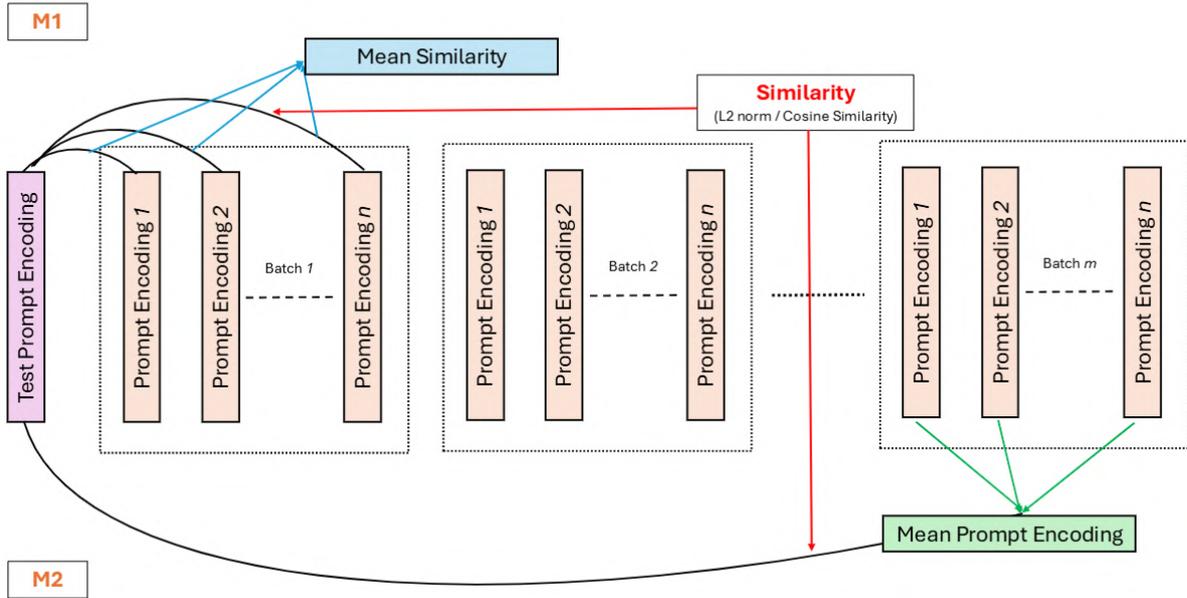


Figure 4: Router Approach for TTS

which can take one of five values - `avg_sim_score`, `avg_prompt_embed`, `max_confidence`, `majority_vote` or (summing log probabilities) i.e., `sum_logprobs` (former two belong to router approach and the latter three constitute the ensemble approach).

- For LS, we use [4] and the corresponding JSON object has fields `times` and `learning_rate`.

6. Experiments

6.1. Setup

As described in Section 5, the base LLM used is Qwen2.5-0.5B-Instruct, domain is common-sense-reasoning and evaluation datasets are ARC-c, BoolQ, HellaSwag and ARC-e. Baselines used include quite recent works such as DnD [28], Test-Time Learning (TTL) [3], Decoupled and Orthogonal Merging (DOM)⁹ [62] and average of task-specific training LoRA's [34]. On one extreme, TTL uses the entire unlabeled corpus of the training LoRA's in addition to the 128 unlabeled examples from the target dataset as seen by SOLAR. On the other extreme, instead of using the unlabeled corpus, DOM merges all the 7 training LoRA's inclusive of the target set.

6.2. Hardware

All experiments were conducted on a high-performance computing node running Ubuntu 22.04.1. The backend processor was EPYC 8434P, which had 48 physical cores (96 logical threads), 256 GB of system RAM and a maximum clock speed of 2.5 GHz. Four NVIDIA RTX A6000 GPUs, each with 48 GB of dedicated VRAM were utilized. Python version used was 3.12.11 and GPU-accelerated tasks were managed using CUDA version 12.4.

⁹DOM is a data-free framework for LoRA merging. It separates parameters into magnitude and direction components and merges them independently, thereby reducing the impact of magnitude differences on the directional alignment of the merged models, thus helping in preserving task information. It also uses a data-free, layer-wise gradient descent method with orthogonal constraints to mitigate interference during the merging of direction components. For evaluation on a target dataset, LoRA's of remaining datasets are merged and used.

Algorithm 1 Sequential Multi-Level RL Loop for Adaptation Strategy (AS) Generation of SOLAR

1: **Input:** Base LM_θ , dataset context C , evaluation metric τ , initial knowledge base K
2: **Init:** Low-rank adapter generator G , sampled adapters $S \leftarrow \text{Sampler}(C, G)$
3: **Level I (Single-edit self-training):**
4: **for** iteration $t = 1, \dots, T_1$ **do**
5: Propose single-edit AS from K $AS \sim LM_\theta(K, C)$
6: Apply AS and obtain weights $\theta' \leftarrow \text{ApplyStrategy}(\theta, AS, S)$
7: Evaluate $Ans \sim LM_{\theta'}(\cdot | \tau)$
8: Compute reward $r \leftarrow r(Ans, \tau)$
9: **if** $r > \text{threshold}_1$ **then**
10: $\theta \leftarrow \text{RL_Update}(\theta, r, AS)$
11: **end if**
12: **end for**
13: **Level II (Chained/compositional strategies):**
14: **for** iteration $t = 1, \dots, T_2$ **do**
15: Propose chain of edits $AS = [e_1, \dots, e_k], e_i \in K$
16: Sequentially apply chain $\theta_0 \leftarrow \theta; \theta_i \leftarrow \text{ApplyStrategy}(\theta_{i-1}, e_i, S)$
17: Evaluate final weights $Ans \sim LM_{\theta_k}(\cdot | \tau)$
18: Compute reward $r \leftarrow r(Ans, \tau)$
19: **if** $r > \text{threshold}_2$ **then**
20: Add chain to KB $K \leftarrow K \cup \{AS\}$
21: $\theta \leftarrow \text{RL_Update}(\theta, r, AS)$
22: **end if**
23: **end for**
24: **Level III (Open-ended exploration):**
25: **for** iteration $t = 1, \dots, T_3$ **do**
26: Generate unconstrained AS $AS \sim LM_\theta(\cdot | C)$ (novel structure)
27: Validate (syntax/safety); if invalid **continue**
28: Apply AS conservatively (strong meta-reg) $\theta' \leftarrow \text{ApplyStrategy}(\theta, AS, S)$
29: Evaluate and compute reward $Ans \sim LM_{\theta'}(\cdot | \tau); r \leftarrow r(Ans, \tau)$
30: **if** $r > \text{threshold}_3$ **then**
31: $K \leftarrow K \cup \{AS\}; \theta \leftarrow \text{RL_Update}(\theta, r, AS)$
32: **else**
33: Penalize harmful proposals in policy update
34: **end if**
35: **end for**
36: **Return:** Refined parameters θ^* , enriched KB K^*

6.3. Results

The major results of this work are presented in Table 1 wherein we conduct experiments of 5 benchmarks which are in the domain of common-sense reasoning and also on 5 out-of-domain benchmarks namely GSM-MC and MATH-MC¹⁰ to evaluate mathematical reasoning, DivLogicEval [66] for logical reasoning, SocialQA [67] for reasoning about social interactions and CodeMMLU [68] for reasoning about code-related tasks. It can be seen that SOLAR in its initial version itself outperforms the task-specific training LoRA’s, TTL, DOM and even DnD by a significant margin, showcasing the promising potential it is capable of, if further levels of RL training¹¹ is completed as well.

Following were the adaptation strategies identified, which enabled SOLAR to reach the accuracy levels presented,

¹⁰GSM-MC and MATH-MC are multiple choice versions of the standard GSM-8K [63] and MATH [64] datasets. They were selected for two reasons - ease of evaluation and correlation with performance on their subjective counterparts [65].

¹¹This might be quite time-intensives however with current version itself taking around 4 days using 2 A6000 GPU’s. The reason for using only 2 despite 4, is because Qwen family has 14 attention heads and the vllm serves used for improved efficiency in inference requires this number to be divisible by the number of GPU’s which is only possible if either 2 or 7 are available.

Dataset	In-Domain Tasks					Avg.	Out-of-Domain Tasks					Avg.
	ARC-e	ARC-c	BoolQ	HellaSwag	PIQA	In	GSM	MATH	Logic	Social	Code	Out
LoRA	47.4	39.7	14.7	26.3	51.5	35.9	15.6	6.8	20.3	39.5	29.8	22.4
TTL	24.4	24.7	44.4	25.9	51.9	34.3	23.5	19.7	26.2	34.9	29.7	26.8
DOM	56.5	38.9	33.2	28.3	18.8	35.1	17.7	2.6	24.7	51.3	31.6	25.6
DnD	70.9	48.1	51.9	26.5	47.8	49.0	20.8	24.1	21.0	33.5	29.1	25.7
SOLAR	74.7	55.5	58.8	48.3	60.1	59.5	30.3	24.5	25.1	55.0	35.6	34.1
Δ DnD	3.8 \uparrow	7.4 \uparrow	6.9 \uparrow	21.8 \uparrow	12.3 \uparrow	10.4 \uparrow	9.5 \uparrow	0.4 \uparrow	4.1 \uparrow	21.5 \uparrow	6.5 \uparrow	8.4 \uparrow
Δ DOM	18.2 \uparrow	16.6 \uparrow	25.6 \uparrow	20.0 \uparrow	41.3 \uparrow	24.3 \uparrow	12.6 \uparrow	21.9 \uparrow	0.4 \uparrow	3.7 \uparrow	4.0 \uparrow	8.5 \uparrow
Δ TTL	50.3 \uparrow	30.8 \uparrow	14.4 \uparrow	22.4 \uparrow	8.2 \uparrow	25.2 \uparrow	6.8 \uparrow	4.8 \uparrow	1.1 \downarrow	20.1 \uparrow	5.9 \uparrow	7.3 \uparrow
Δ LoRA	27.3 \uparrow	15.8 \uparrow	44.1 \uparrow	22.0 \uparrow	8.6 \uparrow	23.6 \uparrow	14.7 \uparrow	17.7 \uparrow	4.8 \uparrow	15.5 \uparrow	5.8 \uparrow	11.7 \uparrow

Table 1

Accuracy (in %) of SOLAR Level I approach over the baselines TTL (25.2 \uparrow), LoRA (23.6 \uparrow), DOM (24.3 \uparrow), and DnD (10.4 \uparrow) for in-domain tasks, and TTL (7.3 \uparrow), LoRA (11.7 \uparrow), DOM (8.5 \uparrow), and DnD (8.4 \uparrow) for out-of-domain tasks, where values in parentheses denote average Δ (change in accuracy) for **Qwen2.5-0.5B-Instruct**.

- For ARC-e and PIQA, it was TTT family with configuration {"ttl_steps": 25, "learning_rate": 1e-5, "batch_size": 4, "shuffle_data": True}
- For ARC-c and SocialIQA, it was LS family with configuration {"times": 5, "learning_rate": 0.1}
- For BoolQ, GSM-MC and MATH-MC, it was LoRA family with TS-mixing strategy and the configuration was {"lambda": 0.5}
- For HellaSwag, DivLogicEval and CodeMMLU, it was TTS family. Ex: for HellaSwag, the corresponding configuration was {"num_prompt_batches": 20, "method": max_confidence}, indicative of the ensemble approach

6.4. Ablation Study

A primary effect we would like to isolate and study is that of initial prompt batch provided to start the LLM adaptation process using SOLAR. It would be ideal if SOLAR results in similar performance even if a highly representative, diverse and influential prompt batch is used. For this purpose, inspired by [53], we use the following strategy for prompt filtering and selection (see Figure 5).

We first model inter-prompt relations as a directed graph $\mathcal{G} = (\mathbf{V}, \mathbf{E}, \mathbf{P})$, wherein each prompt is encoded as a vector by using Sentence-BERT. Each vertex $v_i \in \mathbf{V}$ denotes a prompt (sample), a directed edge $e(i, j) \in \mathbf{E}$ connects v_i to its neighbor v_j , and weight $p(i, j) \in \mathbf{P}$ is the cosine similarity of their embeddings. For each node v_i , an s_i is computed as shown below so that nodes with higher average similarity make more connections.

$$s_i = \frac{1}{|\mathbf{V}| - 1} \sum_{j \neq i} s(i, j), \quad k_i = \lceil \alpha \cdot s_i \cdot (|\mathbf{V}| - 1) \rceil$$

Samples are then scored by (1) influence and (2) diversity. The influence score $I(v)$ is obtained by a diffusion simulation¹². For this, first initialize an active set $S_{\text{active}} = \{v\}$, then iteratively sample an active node u and attempt to activate each neighbor $w \in N_1(u)$ with probability $p(u, w)$. Newly activated nodes join S_{active} . This process is repeated until no active nodes remain. Let $I(v)$ be the total number of visited nodes. Diversity penalty $D(v)$ measures overlap with already selected nodes:

$$D(v) = - \sum_{i=1}^k \beta^i |N_i(v) \cap S_{\text{selected}}|, \quad f_{\mathcal{G}}(v) = I(v) + \gamma D(v)$$

¹²The simulation is run 20 times and is then averaged to obtain the final value.

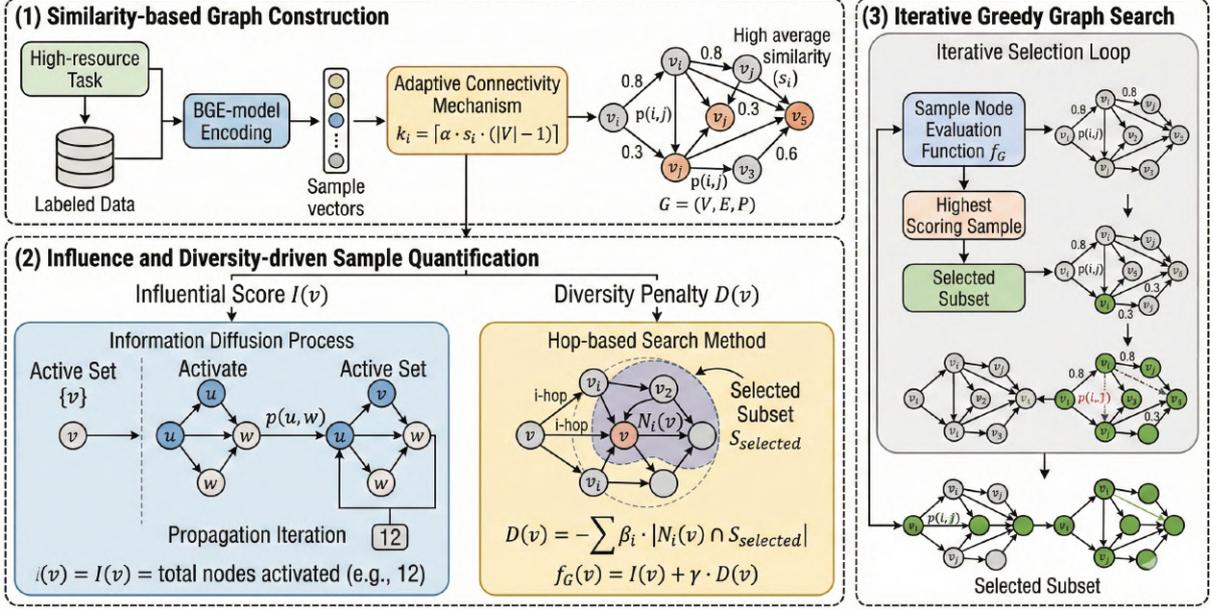


Figure 5: Details of the Prompt Selection Strategy used in Ablation Study

Finally, greedy graph search is done to select the final prompt subset S . For this, start with $S = \emptyset$ and at each round pick

$$v^* = \arg \max_{v \notin S} f_G(v),$$

v^* is then added to S and diversity penalties only for neighbors of v^* are updated¹³. This process continues until $|S|$ reaches the target size which in our case is 128.

Fortunately, the influence of the initial prompt batch was **marginal** (with just a **0.3%** improvement in accuracy when averaged across all evaluation datasets), indicating that SOLAR can efficiently adapt LLMs to unseen datasets without the requirement of high-quality or manually curated dataset. Only a handful of unlabeled prompt instances which are merely indicative of the task suffice.

7. Conclusion

In this paper, we introduce SOLAR which is a novel paradigm for Streaming and Continual Learning by empowering LLMs to autonomously discover and retain parameter-level adaptation strategies. By bridging the gap between rapid test-time adaptation (plasticity) and long-term meta-knowledge retention (stability), SOLAR addresses the core challenges of deploying agents in non-stationary environments. While currently reliant on a seed knowledge base, the framework lays the groundwork for fully autonomous, self-evolving systems capable of navigating the open-ended drifts of the real world. Another key tradeoff is that of real-time adaptation versus computation. While SOLAR’s training phase is compute-intensive, the inference-time application of learned strategies is rapid. By pre-compiling complex adaptation routines into the knowledge base, SOLAR shifts the computational burden from the streaming phase to the offline meta-learning phase. This allows the agent to react to concept drift in near real-time by simply retrieving and applying a cached strategy, rather than performing expensive gradient descent from scratch every time.

¹³Note that the influence scores are precomputed.

8. Acknowledgments

The authors would like to thank Professor Sashikumaar Ganesan, from the Department of Computational and Data Science at Indian Institute of Science, Bangalore for feedback and additional compute resources required to execute this project.

Declaration on Generative AI

During the preparation of this work, the authors used Large Language Models (GPT-5.2, Claude Opus 4.5 and Gemini-3) as a writing assistant tool for drafting content, to generate literature review, for abstract drafting, to paraphrase and reword, to improve writing style, for grammar and spelling check as well as to generate the images used in the paper. The process was interactive. After writing the core content, the authors used LLMs with specific prompts to refine the text. These prompts included requests to “check for grammatical errors,” “rephrase this sentence for clarity,” “make this paragraph more concise,” or “suggest alternative phrasing to improve flow.” The LLMs were not used to generate any scientific ideas, experimental results, data analysis or other core intellectual contributions of the paper. After using these tool(s)/service(s), the authors reviewed and edited the content as needed and take full responsibility for the publication’s content.

References

- [1] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, D. Amodei, Scaling laws for neural language models, arXiv preprint arXiv:2001.08361 (2020).
- [2] W. Wen, C. Wu, Y. Wang, Y. Chen, H. Li, Learning structured sparsity in deep neural networks, Advances in neural information processing systems 29 (2016).
- [3] J. Hu, Z. Zhang, G. Chen, X. Wen, C. Shuai, W. Luo, B. Xiao, Y. Li, M. Tan, Test-time learning for large language models, 2025. URL: <https://arxiv.org/abs/2505.20633>. arXiv: 2505.20633.
- [4] Y. Hu, X. Zhang, X. Fang, Z. Chen, X. Wang, H. Zhang, G. Qi, Slot: Sample-specific language model optimization at test-time, 2025. URL: <https://arxiv.org/abs/2505.12392>. arXiv: 2505.12392.
- [5] Y. Zuo, K. Zhang, L. Sheng, S. Qu, G. Cui, X. Zhu, H. Li, Y. Zhang, X. Long, E. Hua, B. Qi, Y. Sun, Z. Ma, L. Yuan, N. Ding, B. Zhou, Ttrl: Test-time reinforcement learning, 2025. URL: <https://arxiv.org/abs/2504.16084>. arXiv: 2504.16084.
- [6] M. M. Moradi, H. Amer, S. Mudur, W. Zhang, Y. Liu, W. Ahmed, Continuous self-improvement of large language models by test-time training with verifier-driven sample selection, 2025. URL: <https://arxiv.org/abs/2505.19475>. arXiv: 2505.19475.
- [7] H. Lee, S. Oh, J. Kim, J. Shin, J. Tack, Revise: Learning to refine at test-time via intrinsic self-verification, 2025. URL: <https://arxiv.org/abs/2502.14565>. arXiv: 2502.14565.
- [8] J. Hübötter, L. Diaz-Bone, I. Hakimi, A. Krause, M. Hardt, Learning on the job: Test-time curricula for targeted reinforcement learning, 2025. URL: <https://arxiv.org/abs/2510.04786>. arXiv: 2510.04786.
- [9] R. Bertolissi, J. Hübötter, I. Hakimi, A. Krause, Local mixtures of experts: Essentially free test-time training via model merging, 2025. URL: <https://arxiv.org/abs/2505.14136>. arXiv: 2505.14136.
- [10] Z. Yang, N. Band, S. Li, E. Candès, T. Hashimoto, Synthetic continued pretraining, 2024. URL: <https://arxiv.org/abs/2409.07431>. arXiv: 2409.07431.
- [11] Y. Wang, X. Liu, X. Chen, S. O’Brien, J. Wu, J. McAuley, Self-updatable large language models by integrating context into model parameters, 2025. URL: <https://arxiv.org/abs/2410.00487>. arXiv: 2410.00487.
- [12] R. Wang, P. Ping, Z. Guo, X. Zhang, Q. Shi, L. Zhou, T. Ji, Loki: Low-damage knowledge implanting of large language models, 2025. URL: <https://arxiv.org/abs/2505.22120>. arXiv: 2505.22120.
- [13] C. F. Park, Z. Zhang, H. Tanaka, *New News*: System-2 fine-tuning for robust integration of new knowledge, 2025. URL: <https://arxiv.org/abs/2505.01812>. arXiv: 2505.01812.

- [14] L. Chen, M. Prabhudesai, K. Fragkiadaki, H. Liu, D. Pathak, Self-questioning language models, arXiv preprint arXiv:2508.03682 (2025).
- [15] C. Huang, W. Yu, X. Wang, H. Zhang, Z. Li, R. Li, J. Huang, H. Mi, D. Yu, R-zero: Self-evolving reasoning llm from zero data, 2025. URL: <https://arxiv.org/abs/2508.05004>. arXiv: 2508.05004.
- [16] E. C. Acikgoz, C. Qian, H. Ji, D. Hakkani-Tür, G. Tur, Self-improving llm agents at test-time, 2025. URL: <https://arxiv.org/abs/2510.07841>. arXiv: 2510.07841.
- [17] J.-C. Pang, P. Wang, K. Li, X.-H. Chen, J. Xu, Z. Zhang, Y. Yu, Language model self-improvement by reinforcement learning contemplation, 2023. URL: <https://arxiv.org/abs/2305.14483>. arXiv: 2305.14483.
- [18] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, J. Steinhardt, Measuring mathematical problem solving with the math dataset, 2021. URL: <https://arxiv.org/abs/2103.03874>. arXiv: 2103.03874.
- [19] A. Zweiger, J. Pari, H. Guo, E. Akyürek, Y. Kim, P. Agrawal, Self-adapting language models, 2025. URL: <https://arxiv.org/abs/2506.10943>. arXiv: 2506.10943.
- [20] M. Li, J. Lin, X. Zhao, W. Lu, P. Zhao, S. Wermter, D. Wang, Curriculum-rlaif: Curriculum alignment with reinforcement learning from ai feedback, 2025. URL: <https://arxiv.org/abs/2505.20075>. arXiv: 2505.20075.
- [21] W. Yuan, R. Y. Pang, K. Cho, X. Li, S. Sukhbaatar, J. Xu, J. Weston, Self-rewarding language models, 2025. URL: <https://arxiv.org/abs/2401.10020>. arXiv: 2401.10020.
- [22] H. Zhou, Y. Chen, S. Guo, X. Yan, K. H. Lee, Z. Wang, K. Y. Lee, G. Zhang, K. Shao, L. Yang, J. Wang, Memento: Fine-tuning llm agents without fine-tuning llms, 2025. URL: <https://arxiv.org/abs/2508.16153>. arXiv: 2508.16153.
- [23] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, S. Levine, Meta-reinforcement learning of structured exploration strategies, 2018. URL: <https://arxiv.org/abs/1802.07245>. arXiv: 1802.07245.
- [24] K. Irie, I. Schlag, R. Csordás, J. Schmidhuber, A modern self-referential weight matrix that learns to modify itself, 2022. URL: <https://arxiv.org/abs/2202.05780>. arXiv: 2202.05780.
- [25] Z. Tao, T.-E. Lin, X. Chen, H. Li, Y. Wu, Y. Li, Z. Jin, F. Huang, D. Tao, J. Zhou, A survey on self-evolution of large language models, 2024. URL: <https://arxiv.org/abs/2404.14387>. arXiv: 2404.14387.
- [26] H. ang Gao, J. Geng, W. Hua, M. Hu, X. Juan, H. Liu, S. Liu, J. Qiu, X. Qi, Y. Wu, H. Wang, H. Xiao, Y. Zhou, S. Zhang, J. Zhang, J. Xiang, Y. Fang, Q. Zhao, D. Liu, Q. Ren, C. Qian, Z. Wang, M. Hu, H. Wang, Q. Wu, H. Ji, M. Wang, A survey of self-evolving agents: On path to artificial super intelligence, 2025. URL: <https://arxiv.org/abs/2507.21046>. arXiv: 2507.21046.
- [27] K. Wang, D. Tang, W. Zhao, K. Schürholt, Z. Wang, Y. You, Recurrent diffusion for large-scale parameter generation, arXiv preprint arXiv:2501.11587 (2025).
- [28] Z. Liang, D. Tang, Y. Zhou, X. Zhao, M. Shi, W. Zhao, Z. Li, P. Wang, K. Schürholt, D. Borth, et al., Drag-and-drop llms: Zero-shot prompt-to-weights, arXiv preprint arXiv:2506.16406 (2025).
- [29] R. Charakorn, E. Cetin, Y. Tang, R. T. Lange, Text-to-lora: Instant transformer adaption, 2025. URL: <https://arxiv.org/abs/2506.06105>. arXiv: 2506.06105.
- [30] R. M. S. Khan, D. Tang, P. Li, K. Wang, T. Chen, Oral: Prompting your large-scale loras via conditional recurrent diffusion, 2025. URL: <https://arxiv.org/abs/2503.24354>. arXiv: 2503.24354.
- [31] X. Jin, K. Wang, D. Tang, W. Zhao, Y. Zhou, J. Tang, Y. You, Conditional lora parameter generation, 2024. URL: <https://arxiv.org/abs/2408.01415>. arXiv: 2408.01415.
- [32] Y. Shao, X. Lin, X. Long, S. Chen, M. Yan, Y. Liu, Z. Yan, A. Ma, H. Tang, J. Guo, Icm-fusion: In-context meta-optimized lora fusion for multi-task adaptation, 2025. URL: <https://arxiv.org/abs/2508.04153>. arXiv: 2508.04153.
- [33] Y. Shao, M. Yan, Y. Liu, S. Chen, W. Chen, X. Long, Z. Yan, L. Li, C. Zhang, N. Sebe, H. Tang, Y. Wang, H. Zhao, M. Wang, J. Guo, In-context meta lora generation, 2025. URL: <https://arxiv.org/abs/2501.17635>. arXiv: 2501.17635.
- [34] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al., Lora: Low-rank adaptation of large language models, in: International Conference on Learning Representations, 2022, p. 3.
- [35] Y. LeCun, A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27, Open

Review 62 (2022) 1–62.

- [36] Y. Liu, Y. Nan, W. Xu, X. Hu, L. Ye, Z. Qin, P. Liu, Alphago moment for model architecture discovery, 2025. URL: <https://arxiv.org/abs/2507.18074>. arXiv:2507.18074.
- [37] C. Lu, S. Holt, C. Fanconi, A. J. Chan, J. Foerster, M. van der Schaar, R. T. Lange, Discovering preference optimization algorithms with and for large language models, 2024. URL: <https://arxiv.org/abs/2406.08414>. arXiv:2406.08414.
- [38] C. Huang, W. Yu, X. Wang, H. Zhang, Z. Li, R. Li, J. Huang, H. Mi, D. Yu, R-zero: Self-evolving reasoning llm from zero data, arXiv preprint arXiv:2508.05004 (2025).
- [39] N. Reimers, I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, arXiv preprint arXiv:1908.10084 (2019).
- [40] D. Kunin, J. Sagastuy-Brena, S. Ganguli, D. L. Yamins, H. Tanaka, Neural mechanics: Symmetry and broken conservation laws in deep learning dynamics, arXiv preprint arXiv:2012.04728 (2020).
- [41] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16x16 words: Transformers for image recognition at scale, arXiv preprint arXiv:2010.11929 (2020).
- [42] Qwen, :, A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei, H. Lin, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Lin, K. Dang, K. Lu, K. Bao, K. Yang, L. Yu, M. Li, M. Xue, P. Zhang, Q. Zhu, R. Men, R. Lin, T. Li, T. Tang, T. Xia, X. Ren, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Wan, Y. Liu, Z. Cui, Z. Zhang, Z. Qiu, Qwen2.5 technical report, 2025. URL: <https://arxiv.org/abs/2412.15115>. arXiv:2412.15115.
- [43] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, Y. Choi, Hellaswag: Can a machine really finish your sentence?, 2019. URL: <https://arxiv.org/abs/1905.07830>. arXiv:1905.07830.
- [44] C. Clark, K. Lee, M.-W. Chang, T. Kwiatkowski, M. Collins, K. Toutanova, Boolq: Exploring the surprising difficulty of natural yes/no questions, 2019. URL: <https://arxiv.org/abs/1905.10044>. arXiv:1905.10044.
- [45] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, O. Tafjord, Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL: <https://arxiv.org/abs/1803.05457>. arXiv:1803.05457.
- [46] T. Mihaylov, P. Clark, T. Khot, A. Sabharwal, Can a suit of armor conduct electricity? a new dataset for open book question answering, arXiv preprint arXiv:1809.02789 (2018).
- [47] Y. Bisk, R. Zellers, R. L. Bras, J. Gao, Y. Choi, Piqa: Reasoning about physical commonsense in natural language, 2019. URL: <https://arxiv.org/abs/1911.11641>. arXiv:1911.11641.
- [48] K. Sakaguchi, R. L. Bras, C. Bhagavatula, Y. Choi, Winogrande: An adversarial winograd schema challenge at scale, 2019. URL: <https://arxiv.org/abs/1907.10641>. arXiv:1907.10641.
- [49] L. A. Agrawal, S. Tan, D. Soylu, N. Ziemis, R. Khare, K. Opsahl-Ong, A. Singhvi, H. Shandilya, M. J. Ryan, M. Jiang, C. Potts, K. Sen, A. G. Dimakis, I. Stoica, D. Klein, M. Zaharia, O. Khattab, Gepa: Reflective prompt evolution can outperform reinforcement learning, 2025. URL: <https://arxiv.org/abs/2507.19457>. arXiv:2507.19457.
- [50] J. Li, X. Dong, Y. Liu, Z. Yang, Q. Wang, X. Wang, S. Zhu, Z. Jia, Z. Zheng, Reflectevo: Improving meta introspection of small llms by learning self-reflection, 2025. URL: <https://arxiv.org/abs/2505.16475>. arXiv:2505.16475.
- [51] L. Liu, C. Zhang, L. Wu, C. Zhao, Z. Hu, M. He, J. Fan, Instruct-of-reflection: Enhancing large language models iterative reflection capabilities via dynamic-meta instruction, 2025. URL: <https://arxiv.org/abs/2503.00902>. arXiv:2503.00902.
- [52] M. Yuksekgonul, F. Bianchi, J. Boen, S. Liu, Z. Huang, C. Guestrin, J. Zou, Textgrad: Automatic "differentiation" via text, 2024. URL: <https://arxiv.org/abs/2406.07496>. arXiv:2406.07496.
- [53] X. Tang, Z. Lv, X. Cheng, J. Li, W. X. Zhao, Z. Wen, Z. Zhang, J. Zhou, Enhancing cross-task transfer of large language models via activation steering, 2025. URL: <https://arxiv.org/abs/2507.13236>. arXiv:2507.13236.
- [54] T. Wu, J. Wang, Z. Zhao, N. Wong, Mixture-of-subspaces in low-rank adaptation, 2025. URL: <https://arxiv.org/abs/2406.11909>. arXiv:2406.11909.
- [55] R. Wang, K. Dvijotham, I. R. Manchester, Norm-bounded low-rank adaptation, 2025. URL: <https://arxiv.org/abs/2507.13236>.

- //arxiv.org/abs/2501.19050. arXiv:2501.19050.
- [56] Z. Zhao, T. Shen, D. Zhu, Z. Li, J. Su, X. Wang, K. Kuang, F. Wu, Merging loras like playing lego: Pushing the modularity of lora to extremes through rank-wise clustering, 2024. URL: <https://arxiv.org/abs/2409.16167>. arXiv:2409.16167.
 - [57] L. Chen, M. Prabhudesai, K. Fragkiadaki, H. Liu, D. Pathak, Self-questioning language models, 2025. URL: <https://arxiv.org/abs/2508.03682>. arXiv:2508.03682.
 - [58] G. Zhang, F. Meng, G. Wan, Z. Li, K. Wang, Z. Yin, L. Bai, S. Yan, Latentevolve: Self-evolving test-time scaling in latent space, 2025. URL: <https://arxiv.org/abs/2509.24771>. arXiv:2509.24771.
 - [59] A. Karan, Y. Du, Reasoning with sampling: Your base model is smarter than you think, 2025. URL: <https://arxiv.org/abs/2510.14901>. arXiv:2510.14901.
 - [60] Z. Wang, D. Ma, X. Huang, D. Cai, T. Lan, J. Xu, H. Mi, X. Tang, Y. Wang, The end of manual decoding: Towards truly end-to-end language models, 2025. URL: <https://arxiv.org/abs/2510.26697>. arXiv:2510.26697.
 - [61] A. Singh, J. D. Co-Reyes, R. Agarwal, A. Anand, P. Patil, X. Garcia, P. J. Liu, J. Harrison, J. Lee, K. Xu, A. Parisi, A. Kumar, A. Alemi, A. Rizkowsky, A. Nova, B. Adlam, B. Bohnet, G. Elsayed, H. Sedghi, I. Mordatch, I. Simpson, I. Gur, J. Snoek, J. Pennington, J. Hron, K. Kenealy, K. Swersky, K. Mahajan, L. Culp, L. Xiao, M. L. Bileschi, N. Constant, R. Novak, R. Liu, T. Warkentin, Y. Qian, Y. Bansal, E. Dyer, B. Neyshabur, J. Sohl-Dickstein, N. Fiedel, Beyond human data: Scaling self-training for problem-solving with language models, 2024. URL: <https://arxiv.org/abs/2312.06585>. arXiv:2312.06585.
 - [62] S. Zheng, H. Wang, C. Huang, X. Wang, T. Chen, J. Fan, S. Hu, P. Ye, Decouple and orthogonalize: A data-free framework for lora merging, 2025. URL: <https://arxiv.org/abs/2505.15875>. arXiv:2505.15875.
 - [63] K. Cobbe, V. Kosaraju, M. Bavarian, M. Chen, H. Jun, L. Kaiser, M. Plappert, J. Tworek, J. Hilton, R. Nakano, et al., Training verifiers to solve math word problems, arXiv preprint arXiv:2110.14168 (2021).
 - [64] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, J. Steinhardt, Measuring mathematical problem solving with the math dataset, arXiv preprint arXiv:2103.03874 (2021).
 - [65] Z. Zhang, Z. Jiang, L. Xu, H. Hao, R. Wang, Multiple-choice questions are efficient and robust llm evaluators, arXiv preprint arXiv:2405.11966 (2024).
 - [66] T. T. Chung, L. Liu, M. Yu, D.-Y. Yeung, Divlogiceval: A framework for benchmarking logical reasoning evaluation in large language models, arXiv preprint arXiv:2509.15587 (2025).
 - [67] M. Sap, H. Rashkin, D. Chen, R. LeBras, Y. Choi, Socialliqa: Commonsense reasoning about social interactions, arXiv preprint arXiv:1904.09728 (2019).
 - [68] D. N. Manh, T. P. Chau, N. Le Hai, T. T. Doan, N. V. Nguyen, Q. Pham, N. D. Bui, Codemmlu: A multi-task benchmark for assessing code understanding capabilities of codellms, CoRR (2024).