

# Model Whisper: Steering Vectors Unlock Large Language Models’ Potential in Test-time

Xinyue Kang<sup>1</sup>, Diwei Shi<sup>2,1</sup>, Li Chen<sup>1\*</sup>

<sup>1</sup>Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>2</sup>Research Institute of Tsinghua University in Shenzhen

kangxy24@mails.tsinghua.edu.cn, sdw18701584492@163.com, lic87766@gmail.com

## Abstract

It is a critical challenge to efficiently unlock the powerful reasoning potential of Large Language Models (LLMs) for specific tasks or new distributions. Existing test-time adaptation methods often require tuning model parameters, which is not only computationally expensive but also risks degrading the model’s pre-existing abilities. To address this, we introduce a lightweight component, *Test-Time Steering Vectors* (TTSV), which is prepended to the input while keeping the LLM’s parameters entirely frozen. By optimizing the TTSV on test data to minimize the model’s output entropy, we steer the model towards an internal state of higher confidence, activating its inherent abilities most relevant to the current task. TTSV is both lightweight and highly efficient to optimize, making it a true plug-and-play enhancement. Extensive experiments validate our approach’s effectiveness on both base models and reasoning-enhanced models. For instance, on the MATH500 task, TTSV achieves a 45.88% relative performance gain on the Qwen2.5-Math-7B model and a 16.22% relative gain on the Qwen3-4B model. Furthermore, our approach exhibits robust generalization, with its steering vectors proving highly transferable across diverse tasks.

**Code** — <https://github.com/kkkkxy/TTSV>

## 1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities (Jaech et al. 2024; Guo et al. 2025), evolving into powerful general-purpose problem solvers. Trained on massive corpora of text and code, these models internalize a wealth of world knowledge and complex reasoning patterns (Wang et al. 2024). However, this immense potential often remains latent. When confronted with a specific task or a novel data distribution, a model may fail to automatically or optimally deploy its intrinsic capabilities, even though they exist (Xu et al. 2025). This gap between the model’s latent potential and its actual performance presents a key challenge to fully realizing its value.

Prevailing methods for guiding large models toward desired behaviors primarily involve post-training, such as fine-tuning the entire model via Reinforcement Learning from

Human Feedback (RLHF) (Bai et al. 2022) or Reinforcement Learning with Verifiable Rewards (RLVR) (Guo et al. 2025). These approaches, though effective, demand substantial computational resources and extensive, costly human-annotated data. Test-time adaptation (TTA) offers a more lightweight alternative by dynamically adapting the model to unlabeled test data. Yet, despite their proven efficacy, prevailing TTA strategies still fundamentally depend on parameter modification, often involving the update of all model parameters at test time, as seen in approaches like Test Time Reinforcement Learning (TTRL) (Zuo et al. 2025) and One-shot Entropy Minimization (EM) (Gao et al. 2025). This dependency on parameter modification gives rise to two critical challenges: 1) the considerable computational expense of gradient calculation and backpropagation, and 2) the inherent risk of catastrophic forgetting, whereby adapting to a new domain might unintentionally erase the model’s invaluable, pretrained general abilities (Liu et al. 2021).

The parameter-efficient TTA method SLOT (Hu et al. 2025) offers an alternative by introducing a learnable parameter vector to the final hidden layer before the output head, acting as a form of output-level calibration. However, due to its limited influence on the final computational step, the entire upstream multi-layer reasoning process remains unguided, restricting its performance.

The question then becomes: how can we guide the LLM’s entire reasoning process from its inception, activating its latent potential, while still keeping its core parameters untouched? This requires a paradigm shift from model modification to holistic model activation—unlocking rather than reshaping its inherent capabilities. While Prompt Engineering (Sahoo et al. 2024) also intervenes at the input stage and preserves model parameters, it essentially “shouts” at the model through discrete, human-readable text. The expressive power of discrete vocabulary consequently limits its effectiveness, whereas the model’s native representation space is continuous and capable of encoding denser, more fine-grained information.

Accordingly, we propose Model Whispering. Instead of “shouting”, we learn to directly “whisper” a cryptic instruction into the model’s continuous embedding space. This signal, though opaque to humans, is information-rich and designed to precisely guide the model’s internal state into the optimal configuration for solving the current task.

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

We realize the “Model Whispering” concept as a lightweight component named the Test-Time Steering Vectors (TTSV). The TTSV is a set of short, optimizable vectors that are directly prepended to the front of the input text’s embedding sequence. This placement is deliberate. By intervening at the very beginning, TTSV guides the entire multi-layer reasoning process, a fundamental distinction from output-level adjustments like SLOT. Specifically, as the prepended vectors participate in the computation within the attention modules across all layers, the output of each module is effectively a linear combination of its original output and a learned bias vector (Petrov, Torr, and Bibi 2023). This bias vector serves to guide the model’s computational trajectory into an optimal configuration, all while preserving the relative attention patterns among the original input tokens.

To optimize TTSV, our approach is guided by a core premise: a model that successfully activates its relevant capabilities will exhibit greater certainty in its predictions. Since output entropy directly measures this predictive uncertainty, we define the optimization goal for the TTSV as minimizing the output entropy. This objective compels the model to converge towards a more definitive and confident judgment, in turn achieving the precise activation of its latent potential. The entire framework is remarkably efficient and non-invasive, requiring just a few gradient updates to the minuscule TTSV and rendering it a genuine plug-and-play enhancement.

We conduct experiments on a wide range of LLMs and benchmarks to validate the effectiveness of our approach. For instance, on the MATH500 dataset, TTSV achieves a 45.88% relative performance gain on the Qwen2.5-Math-7B model (lifting its score from 51.00% to 74.40%), and a 16.22% relative gain on the Qwen3-4B reasoning model (lifting its score from 51.80% to 60.20%). Furthermore, the approach demonstrates robust generalization capabilities. The TTSV optimized on MATH500, for example, boosts performance on the unseen AMC23 dataset by a substantial 20.00%. This interesting phenomenon highlights a key property of our method: a TTSV trained on one task can be effectively transferred to another, consistently unlocking significant performance improvements without re-optimization.

Our contributions can be summarized as follows.

- We propose a novel, light-weight paradigm for TTA. By optimizing an external, plug-and-play component (TTSV) prepended to the input, our method steers the model’s entire reasoning process from its inception. This characteristic fundamentally distinguishes our approach from methods that either modify model weights or only perform output-level calibration, offering a new, non-invasive, and more holistic path for guiding LLMs.
- We demonstrate that TTSV can be effectively optimized using a simple, label-free objective: entropy minimization. This process successfully activates the model’s latent abilities without needing complex reward functions or access to internal gradients.
- Extensive experiments demonstrate our method’s high effectiveness, strong generalization, and preservation of inherent capabilities.

## 2 Related Work

### 2.1 Test-Time Adaptation (TTA)

TTA aims to improve a pre-trained model’s performance by adapting it to unlabeled test instances, thereby addressing distribution shifts encountered during deployment (Liang, He, and Tan 2025). Pioneered in computer vision, foundational TTA approaches often focused on minimizing prediction uncertainty. A prominent example is TENT (Wang et al. 2020), which minimizes the entropy of predictions to adapt batch normalization layers. Other methods, like Test-Time Training (TTT) (Sun et al. 2020), employ self-supervised tasks such as rotation prediction for adaptation.

More recently, the principles of TTA have been applied to LLMs. For instance, TTRL (Zuo et al. 2025) generates pseudo-labels for test data via a majority vote and then performs full-parameter test-time training using reinforcement learning. Another approach, SLOT (Hu et al. 2025), optimizes a parameter vector added to the model’s final hidden layer by performing next-token prediction on the input prompt itself. While these methods have demonstrated efficacy, they face fundamental challenges. The risk of catastrophic forgetting during parameter modification (Liu et al. 2021; Zhao et al. 2023), and the question of how to more deeply guide the model to adapt to a new distribution rather than merely calibrating its output, remain open problems. Consequently, designing lightweight, efficient, and safe TTA methods specifically tailored for LLMs remains a key research direction.

### 2.2 Parameter-Efficient Fine-Tuning (PEFT)

PEFT methods specialize pretrained models for downstream applications by optimizing a minimal subset of parameters during a dedicated fine-tuning phase, while leaving the vast majority of the original LLM weights frozen. Notable examples of these techniques include Adapter Tuning (Houlsby et al. 2019), LoRA (Hu et al. 2022), Prompt Tuning (Lester, Al-Rfou, and Constant 2021), and Prefix-Tuning (Li and Liang 2021).

A common principle unites these methods: they are employed during a training phase, often in a supervised setting, to embed new and task-specific knowledge into the trainable components. The goal is to effectively teach the model a new skill or specialize its behavior for a particular downstream task. In stark contrast to this training-centric philosophy, our approach has a different objective: to activate and steer the model’s pre-existing capabilities at test-time. Operating on unlabeled data, our method focuses on dynamically unlocking what the model already knows, rather than instilling new, task-specific information.

### 2.3 Entropy-based Optimization

The concept of entropy, as a measure of uncertainty, serves as a powerful, dual-purpose objective in machine learning. On one hand, entropy maximization is often employed as a regularization technique, particularly in reinforcement learning (RL), where the expected reward objective is augmented with a term designed to increase the policy’s entropy (Ziebart et al. 2008). The goal is to promote exploration by

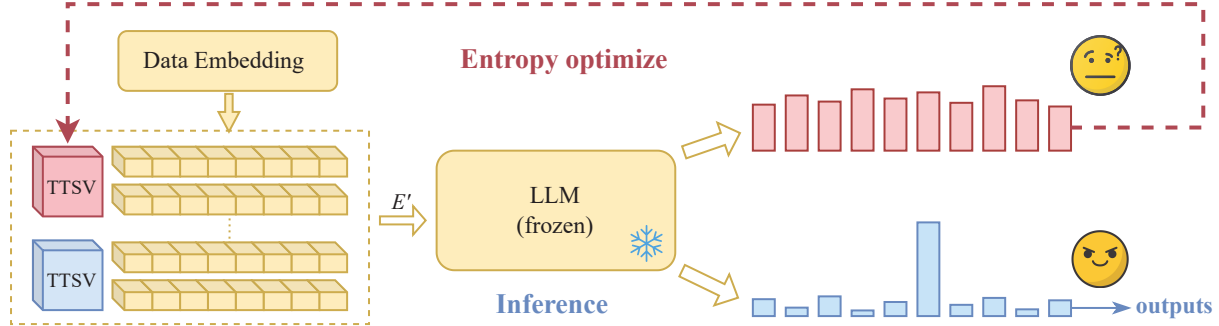


Figure 1: The overall framework of our proposed Model Whisper method. The **upper path (red)** illustrates the optimization phase, where the TTSV is iteratively updated by minimizing the model’s output entropy. The **lower path (blue)** shows the inference phase, where the fixed, optimized TTSV guides the model to make low-entropy, high-confidence predictions.

encouraging uncertainty in the model’s outputs or actions (Haarnoja et al. 2018; Eysenbach and Levine 2021; Wang et al. 2025). This approach is well-suited for scenarios requiring extensive exploration, such as learning a policy from scratch. However, for a capable LLM that already possesses a strong base of knowledge and reasoning patterns, forcing such randomized exploration can be counterproductive. It may destabilize the model’s coherent thought processes, leading to less reliable and structured outputs.

On the other hand, entropy minimization has a long history as a principle for encouraging model confidence and exploiting learned knowledge. It is a well-established regularization technique in semi-supervised learning (Grandvalet and Bengio 2004; Berthelot et al. 2019; Niu et al. 2014), domain adaptation (Roy et al. 2019; Saito et al. 2019), and few-shot learning (Dhillon et al. 2019; Kim, Seo, and Han 2022). More recently, a number of LLM-related works have also leveraged entropy minimization for optimization (Gao et al. 2025; Agarwal et al. 2025; Prabhudesai et al. 2025).

Since TTA is applied to an LLM with substantial latent abilities, the primary challenge is not to discover new strategies but to reliably activate the correct ones. Therefore, we argue that minimizing entropy, thereby reducing uncertainty and steering the model towards a state of high confidence, is a direct and effective strategy for unlocking the task-relevant knowledge within a pre-trained LLM.

### 3 Method

Herein, we detail our proposed Model Whisper, which is designed to guide the LLM by optimizing a set of external, lightweight vectors, TTSV, on unlabeled test data, enhancing LLM’s performance on specific tasks without modifying any of the model’s internal parameters. Figure 1 presents an overview of the entire framework.

#### 3.1 Formulation

Given a pre-trained LLM, denoted as  $M(\cdot; \theta)$ , with its parameters  $\theta$  entirely frozen, our objective is to enhance its performance on an unlabeled test dataset  $\mathcal{D}_{\text{test}} = \{X_1, X_2, \dots, X_N\}$ .

To this end, we introduce the TTSV, learnable continuous vectors denoted as  $\mathbf{V}_{\text{steer}} \in \mathbb{R}^{L \times d}$ . Here,  $L$  represents the sequence length of the steering vectors (a small hyperparameter, e.g.,  $L = 20$ ), and  $d$  is the model’s embedding dimension. The operational mechanism of TTSV is straightforward: for any input text  $X$ , we first retrieve its embedding sequence,  $E(X)$ . The optimized tensor,  $\mathbf{V}_{\text{steer}}$ , is then prepended to this sequence, creating an augmented input  $E' = [\mathbf{V}_{\text{steer}}; E(X)]$ . This augmented sequence is subsequently fed into the frozen LLM  $M$  to generate the final output. In this manner,  $\mathbf{V}_{\text{steer}}$  modulates the model’s initial computational state, thereby steering its subsequent reasoning process without altering its intrinsic parameters.

#### 3.2 Optimization

Our optimization strategy is based on a core premise: a model proficient in a specific task should exhibit high certainty in its predictions, which corresponds to low output entropy. By inverting this relationship, we adopt entropy minimization as our optimization objective. We aim to actively steer the model towards this state of high certainty by reducing its output entropy, thereby activating its inherent, task-relevant reasoning abilities.

When the model processes the augmented input embedding  $E' = [\mathbf{V}_{\text{steer}}; E(X)]$ , it autoregressively generates a token sequence  $Y = (y_1, y_2, \dots, y_T)$ . For the  $t$ -th step of this generation, the conditional entropy  $H_t$  is defined as:

$$H_t = - \sum_{v \in \mathcal{V}} p(v|y_{<t}, E') \log p(v|y_{<t}, E'). \quad (1)$$

To focus the optimization on the model’s generative capabilities, our final loss function is defined as the average token-level entropy across all non-padding generated tokens within a mini-batch. For a mini-batch of  $B$  samples, the loss function  $\mathcal{L}$  is computed as:

$$\mathcal{L}(\mathbf{V}_{\text{steer}}) = \frac{\sum_{i=1}^B \sum_{t \in \mathcal{I}_i} H_{i,t}}{\sum_{i=1}^B |\mathcal{I}_i|}, \quad (2)$$

where  $\mathcal{I}_i$  denotes the set of all non-padding generated token positions for the  $i$ -th sample, and  $H_{i,t}$  is the conditional

Model	MATH500	Minerva Math	Olympiad	AMC23	Avg.
Qwen2.5-Math-7B	51.00	12.90	16.70	42.50	30.78
+ TTSV	74.40	22.80	29.80	65.00	48.00
$\Delta$ (EM)	+15.00	+7.80	+16.00	+17.50	+14.08
$\Delta$ (TTSV)	+23.40 <sup>↑45.88%</sup>	+9.90 <sup>↑76.74%</sup>	+13.10 <sup>↑78.44%</sup>	+22.50 <sup>↑52.94%</sup>	+17.22 <sup>↑55.95%</sup>
LLaMA-3.1-8B	50.60	17.30	11.60	25.00	26.13
+ TTSV	49.40	17.60	14.20	25.00	26.55
$\Delta$ (EM)	+0.20	+1.90	-1.60	+0.30	+0.20
$\Delta$ (TTSV)	-1.20 <sup>↓2.37%</sup>	+0.30 <sup>↑1.73%</sup>	+2.60 <sup>↑22.41%</sup>	+0.00 <sup>↑0.00%</sup>	+0.42 <sup>↑1.61%</sup>

Table 1: Performance Comparison of TTSV against EM on the Qwen2.5-Math-7B and LLaMA-3.1-8B Models.

entropy for sample  $i$  at time step  $t$ .

This loss function is differentiable with respect to  $\mathbf{V}_{\text{steer}}$ . With backpropagation, we compute and update gradients only for  $\mathbf{V}_{\text{steer}}$ , while the model parameters  $\theta$  remain frozen. This unsupervised entropy minimization process, without relying on any external labels, effectively compels the model to converge to a more deterministic internal configuration, thereby unlocking its potential for the task at hand.

### 3.3 Inference

Once TTSV optimized via the process described above, we obtain fixed steering vectors, denoted as  $\mathbf{V}_{\text{steer}}^*$ , which is tailored to the current test distribution. During the inference phase, these optimized steering vectors are used for all samples in the test set.

For any given instance  $X_i \in \mathcal{D}_{\text{test}}$ , the inference process is straightforward: the fixed steering tensor  $\mathbf{V}_{\text{steer}}^*$  is prepended to its corresponding embedding sequence  $E(X_i)$  to construct the augmented input  $E'_i = [\mathbf{V}_{\text{steer}}^*, E(X_i)]$ . This augmented sequence is then fed into the frozen LLM to generate the final prediction.

This ‘optimize-once, apply-to-all’ paradigm renders TTSV a true plug-and-play component. The overhead at inference time is negligible, limited only to the processing of  $L$  additional vectors at the input layer. Consequently, it enhances model performance with virtually no impact on inference speed.

## 4 Experiment

### 4.1 Experimental Setup

**Models.** To demonstrate the broad applicability of our method, we conducted experiments on both base models and reasoning-enhanced models. For base models, we selected models from both the Qwen and LLaMA series to ensure architectural diversity. Specifically, we used Qwen2.5-Math (Yang et al. 2024) in two different scales (1.5B and 7B) and the LLaMA-3.1-8B-Instruct model (Grattafiori et al. 2024). For the reasoning-enhanced category, we chose the latest Qwen3-4B model (Yang et al. 2025) operating in its ‘thinking mode’ to evaluate our method’s effectiveness on state-of-the-art architectures.

**Benchmarks.** Our evaluation is conducted on a diverse suite of benchmarks to test mathematical and general rea-

soning. For mathematical reasoning, we employ widely-used datasets such as MATH500 (Hendrycks et al. 2021), Minerva Math (Lewkowycz et al. 2022), Olympiad Bench (He et al. 2024), and AMC23. To assess performance on problems of higher difficulty, our evaluation also includes AIME24 (Li et al. 2024), derived from the American Invitational Mathematics Examination 2024, which tests advanced, competition-style mathematical problem-solving. In addition, to assess the generalization of our method beyond the mathematical domain, we utilize GPQA Diamond (Rein et al. 2024), a benchmark comprising graduate-level scientific questions in fields like physics and biology, specifically designed to be resistant to web-search retrieval.

**Baselines.** We benchmark our method against two representative test-time adaptation methods: EM (Entropy Minimization) (Gao et al. 2025) and SLOT (Hu et al. 2025). EM represents a full fine-tuning approach that adjusts all model parameters at test time, while SLOT is a parameter-efficient method that performs sample-specific fine-tuning on the output vector of the model’s last hidden layer. We validate the effectiveness of our method through comparison with these approaches.

**Implementation Details.** We set the length  $L$  of TTSV to 20. The TTSV was optimized for 20 epochs using the AdamW optimizer. We employed a linear learning rate schedule, with the rate initialized at 1e-3 and decaying to 1e-5. Other optimization hyperparameters included a weight decay of 1e-8, an epsilon of 1e-5, and a training batch size of 16. During inference, we used a batch size of 64, and the maximum generation length was set to 3072 tokens. All experiments were conducted on NVIDIA A800 GPUs.

### 4.2 Main Results

**Comparison with Full Parameter TTA (EM).** Table 1 shows the experimental results of TTSV and the full fine-tuning EM on the Qwen2.5-Math-7B model and LLaMA-3.1-8B model. On the Qwen2.5-Math-7B model, TTSV demonstrates significant efficacy, achieving an average absolute improvement of +17.22%, which corresponds to a 55.95% relative gain. This result surpasses the +14.08% average gain from EM. On the LLaMA-3.1-8B model, while the gains from both methods were modest, TTSV still achieved a slightly higher average improvement than EM.

Model	AIME24	MATH500	GPQA	Avg.
Qwen2.5-Math-1.5B	6.67	39.40	27.27	24.45
+ TTSV	6.67	68.40	29.80	34.96
$\Delta$ (SLOT)	+3.33	+0.00	-0.51	+0.94
$\Delta$ (TTSV)	+0.00 $\uparrow$ 0.00%	+29.00 $\uparrow$ 73.60%	+2.53 $\uparrow$ 9.28%	+10.51 $\uparrow$ 42.99%
Qwen2.5-Math-7B	13.33	51.00	28.79	31.04
+ TTSV	23.33	74.40	31.82	43.18
$\Delta$ (SLOT)	+6.67	+1.20	+7.07	+4.98
$\Delta$ (TTSV)	+10.00 $\uparrow$ 75.02%	+23.40 $\uparrow$ 45.88%	+3.03 $\uparrow$ 10.52%	+12.14 $\uparrow$ 39.12%

Table 2: Performance Comparison of TTSV against SLOT on the Qwen2.5-Math-1.5B and Qwen2.5-Math-7B Models.

Model	MATH500	AMC23	Minerva Math	Olympiad	AIME24	GPQA	Avg.
Qwen3-4B	51.80	37.50	24.60	18.10	56.67	41.92	38.43
+ TTSV	60.20	40.00	25.40	30.70	60.00	47.98	44.05
$\Delta$	+8.40 $\uparrow$ 16.22%	+2.50 $\uparrow$ 6.67%	+0.80 $\uparrow$ 3.25%	+12.60 $\uparrow$ 69.61%	+3.33 $\uparrow$ 5.88%	+6.06 $\uparrow$ 14.46%	+5.62 $\uparrow$ 14.62%

Table 3: Performances of TTSV on Reasoning Model.

These results suggest that our approach provides an effective path to enhance model performance while greatly reducing the computational overhead typically associated with full fine-tuning during test-time adaptation.

#### Comparison with Parameter-Efficient TTA (SLOT).

We further analyze our method by comparing it with SLOT, a parameter-efficient TTA approach. While both methods keep core model parameters frozen and operate on a small number of tunable parameters, they differ in their point of intervention. SLOT performs a ‘‘post-computation adjustment’’ by modifying the output features of the model’s final layer, acting akin to an output filter. In contrast, our TTSV implements a ‘‘pre-computation steering’’ mechanism, guiding the model’s entire reasoning process from the outset by prepending optimized vectors to the input.

As shown in Table 2, TTSV achieves superior performance on both the Qwen2.5-Math-1.5B and 7B models.

However, the two methods employ different adaptation strategies. SLOT’s ‘‘sample-specific’’ approach offers flexibility, as it can adapt without requiring the entire test set, albeit at the cost of inference latency due to per-sample optimization. We acknowledge the value of this strategy in certain scenarios. To provide a direct comparison, we also conducted a sample-specific experiment with TTSV on the Qwen-2.5-math-1.5B model. As shown in Table 4, the results demonstrate that even in this setting, TTSV still outperforms SLOT. The superior outcomes under both adaptation strategies suggest that guiding the model’s internal state early in the computational graph may be a more effective way to activate latent abilities than adjusting its final output.

Therefore, based on a comprehensive consideration of performance, stability, and inference efficiency, we adopt the ‘‘optimize-once, apply-to-all’’ paradigm as the primary approach in this work.

Model	AIME24	MATH500	GPQA
Baseline	6.67	39.40	27.27
+ TTSV	10.00	43.40	27.78
$\Delta$ (SLOT)	+3.33	+0.00	-0.51
$\Delta$ (TTSV)	+3.33	+4.00	+0.51

Table 4: Performance Comparison of TTSV (sample-specific) against SLOT on the Qwen2.5-Math-1.5B Model.

**TTSV for Reasoning Model.** To assess the applicability of our method to state-of-the-art LLMs, we further evaluated TTSV on a reasoning-enhanced model. Such models, having been specifically post-trained for structured, complex reasoning, present a new challenge for our approach. The task for TTSV is no longer just to activate a latent ability, but to precisely steer an already sophisticated reasoning mechanism without disrupting its inherent strengths. We selected the latest Qwen3-4B model (thinking mode) for this purpose and compared its performance with and without TTSV. As detailed in Table 3, TTSV demonstrates performance gains across all benchmarks. The effect is particularly notable on Olympiad Bench, where it achieves a substantial relative increase of 69.61%. Across all datasets, TTSV achieves an average relative improvement of 14.62%. These results indicate that TTSV, despite a lightweight external component, can effectively synergize with advanced reasoning models, helping to further focus and refine their reasoning pathways. This confirms that our approach is not only effective for base models but also holds significant relevance and potential for the frontier of LLM development.

### 4.3 Ablation Studies

**Analysis of Training Dynamics.** To investigate the training dynamics of TTSV, we tracked the relationship between the entropy-based loss and task accuracy. As depicted in Fig-

Figure 2, as the training loss steadily decreases, the model’s accuracy consistently rises until both metrics converge to a plateau. Notably, this process is stable and does not exhibit the performance degradation at later training stages observed in EM.




Figure 2: Training dynamics of Qwen2.5-Math-1.5B on MATH500. The plot illustrates the change in training loss (left y-axis, red line) and accuracy (right y-axis, blue line) over 20 epochs.

We attribute this stability to our approach’s fundamental design, which operates on an external component while leaving the model’s internal parameters untouched. In contrast, methods like EM, which directly modify the core weights of the model, are susceptible to overfitting on the test distribution. The entropy minimization objective can push the model into a state of over-confidence, leading to overly deterministic and shallow reasoning (Zhang et al. 2025), where the model learns to be certain but incorrect.

By keeping the LLM’s parameters frozen, TTSV preserves the integrity of the model’s core reasoning abilities and its vast knowledge base. The optimization process is thereby constrained to a search for an optimal activation signal within the model’s existing latent space, rather than a reshaping of the space itself. This inherently guards against overfitting and the risk of catastrophic forgetting, ensuring the model’s foundational capabilities remain intact.

**Impact of TTSV Length** The length of the TTSV,  $L$ , is a key hyperparameter that determines the representational capacity of the steering signal. We evaluated its impact by varying the length, setting  $L \in \{1, 5, 10, 20, 40\}$ . The results, presented in Figure 3, yield two important findings.




Figure 3: The effect of TTSV length on the accuracy of Qwen2.5-Math-1.5B on the MATH500 benchmark.

First, the relationship between TTSV length and performance is non-monotonic. Accuracy improves as  $L$  increases from 1 to 20, with  $L = 20$  yielding the best results, but performance declines at  $L = 40$ . This suggests a trade-off: a short prefix may have insufficient capacity to encode an effective steering signal, while an excessively long prefix increases the number of parameters to be optimized and may introduce noise or unfavorably alter the attention dynamics over the primary task input. Thus, an optimal length exists that balances expressive capacity with optimization stability.

Second, it is noteworthy that a TTSV of length  $L = 1$  still provides a substantial improvement over the baseline. This finding powerfully demonstrates the high information density of the model’s continuous embedding space. It suggests that a single, well-optimized vector is sufficient to convey a potent steering command to the model. The ability to guide a model’s behavior with such a minimal and continuous signal hints at a promising area for further investigation into efficient model control.

**Analysis of Cross-Distribution Generalization** To assess the transferability of the steering signals learned by TTSV, we conducted a cross-distribution generalization experiment, with results presented in Figure 4. A consistent finding from this experiment is that applying an Out-of-Distribution (OOD) TTSV, one trained on a different distribution, always yields substantial performance gains over the baseline. For instance, the TTSV trained on Math500 improved accuracy on AMC23 from 42.5% to 62.5%. This indicates that the learned steering signal captures a significant degree of cross-distribution utility, rather than simply overfitting to the source distribution.

Building on this, we observed an interesting phenomenon where, on some tasks, the OOD application outperformed the In-Distribution (ID) TTSV optimized for the target. A clear example is seen when the TTSV trained on Math500 is applied to the Minerva Math task, recording an accuracy of 26.5%, which surpasses the 22.8% from the ID-TTSV trained directly on Minerva Math. We attribute this to the possibility that different training distributions shape TTSVs to specialize in distinct reasoning capabilities. While an ID-TTSV must find a comprehensive solution for all problems within its distribution, an OOD-TTSV may happen to specialize in activating a core ability that is particularly critical for the target task, thus yielding superior performance.

## 5 Theoretical Analysis

We theoretically elaborate on the core mechanism of our proposed TTSV method. Our central thesis is that by prepending a learnable vector to the model’s input, TTSV introduces an initial bias into the computation. This signal is then propagated and amplified through the deep network, systematically steering the model’s computational trajectory towards a desired target state.

We begin by analyzing the precise effect of TTSV within the first attention layer of the model. Ignoring position encoding, when a learnable vector  $\mathbf{p}$  is introduced, the new output  $t'_i$  at any content position  $i$  can be expressed exactly as a linear combination of the original output  $t_i$  and a bias




Figure 4: Cross-distribution generalization of TTSV. The figure compares the performance of the Qwen2.5-Math-7B model with In-Distribution (ID) and Out-of-Distribution (OOD) TTSV applications across four math datasets.

direction  $\mathbf{b}$ :

$$\mathbf{t}'_i = (1 - \alpha_i)\mathbf{t}_i + \alpha_i\mathbf{b}. \quad (3)$$

The bias direction  $\mathbf{b} = \mathbf{W}_V\mathbf{p}$  is determined solely by the learnable vector  $\mathbf{p}$ , while the bias strength  $\alpha_i$  is the attention weight from position  $i$  to  $\mathbf{p}$ . Eq. 3 demonstrates that TTSV introduces a linear guiding signal in the first layer, shifting the model’s activation state from its original trajectory.

This initial guiding signal propagates through the model’s deep network. The computation at each layer operates on the already-guided trajectory from the previous one, which allows a small initial offset to be significantly amplified after passing through multiple nonlinear transformations. Our experiments also corroborate this trend of “bias amplification”, empirically showing that the magnitude of the activation shift induced by TTSV increases with network depth.

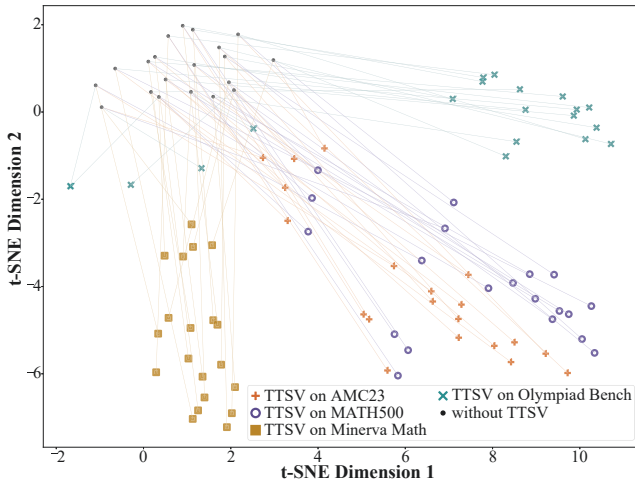


Figure 5: t-SNE visualization of the final layer activations from the Qwen2.5-math-7B model on 20 problems sampled from Minerva Math.

The cumulative effect of this process is manifested in the final layer’s activations, as visualized by the t-SNE plot in Figure 5. In the figure, the baseline activations (black dots)

are systematically pushed towards new, task-specific regions (colored markers) under the influence of TTSV. The vectors connecting the baseline and task-specific points vividly illustrate this aggregate guiding force. This provides strong evidence that TTSV, by injecting a controllable bias early in the computational flow and leveraging the model’s depth to amplify this signal, successfully steers the model’s general-purpose capabilities into the specialized states required for specific tasks. Furthermore, we observe an overlap in the state distributions guided by TTSV vectors trained on different tasks. This observation may explain the strong OOD performance of TTSV reported in Section 4.3, as it suggests the model is guided towards a shared, robust reasoning subspace rather than narrow, task-specific solutions.

## 6 Conclusion

In this paper, we introduced a novel method to enhance the reasoning capabilities of Large Language Models without modifying their internal parameters. By optimizing a small set of input vectors to minimize output entropy, TTSV effectively steers the model towards a state of higher confidence. Our comprehensive experiments and theoretical analyses demonstrate that this lightweight, plug-and-play method is capable of guiding the model’s trajectory towards a target state from the very outset, achieving significant performance gains on both base and reasoning-enhanced models, and exhibiting remarkable generalization across different tasks.

This work illuminates a promising and efficient path, as an alternative to full fine-tuning, to unlock the latent potential of LLMs. Furthermore, this work suggests several potential open questions. First, beyond entropy minimization, the exploration of alternative or combined objective functions remains an important open question. Second, integrating external guidance mechanisms like TTSV with training methods such as reinforcement learning might offer new perspectives on model optimization. Finally, extending the application of this approach from LLMs to Multimodal Large Language Models (MLLMs) also warrants investigation. We hope this work will inspire further research into resource-efficient ways to guide and improve existing models.

## Acknowledgments

This research was supported by National Science Foundation of China-Youth Fund-62401329 (2024-NSFC-62401329) and the China Postdoctoral Science Foundation under Grant Number 2025M772920. We also acknowledge the High Performance Computing Center, Tsinghua University for providing computational resources.

## References

- Agarwal, S.; Zhang, Z.; Yuan, L.; Han, J.; and Peng, H. 2025. The unreasonable effectiveness of entropy minimization in llm reasoning. *arXiv preprint arXiv:2505.15134*.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Berthelot, D.; Carlini, N.; Goodfellow, I.; Papernot, N.; Oliver, A.; and Raffel, C. A. 2019. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32.
- Dhillon, G. S.; Chaudhari, P.; Ravichandran, A.; and Soatto, S. 2019. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*.
- Eysenbach, B.; and Levine, S. 2021. Maximum entropy RL (provably) solves some robust RL problems. *arXiv preprint arXiv:2103.06257*.
- Gao, Z.; Chen, L.; Zhou, J.; and Dai, B. 2025. One-shot Entropy Minimization. *arXiv preprint arXiv:2505.20282*.
- Grandvalet, Y.; and Bengio, Y. 2004. Semi-supervised learning by entropy minimization. *Advances in neural information processing systems*, 17.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. Pmlr.
- He, C.; Luo, R.; Bai, Y.; Hu, S.; Thai, Z. L.; Shen, J.; Hu, J.; Han, X.; Huang, Y.; Zhang, Y.; et al. 2024. Olympiad-bench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International conference on machine learning*, 2790–2799. PMLR.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, Y.; Zhang, X.; Fang, X.; Chen, Z.; Wang, X.; Zhang, H.; and Qi, G. 2025. SLOT: Sample-specific Language Model Optimization at Test-time. *arXiv preprint arXiv:2505.12392*.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Kim, M.; Seo, S.; and Han, B. 2022. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12912–12921.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Lewkowycz, A.; Andreassen, A.; Dohan, D.; Dyer, E.; Michalewski, H.; Ramasesh, V.; Slone, A.; Anil, C.; Schlag, I.; Gutman-Solo, T.; et al. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35: 3843–3857.
- Li, J.; Beeching, E.; Tunstall, L.; Lipkin, B.; Soletskyi, R.; Huang, S.; Rasul, K.; Yu, L.; Jiang, A. Q.; Shen, Z.; et al. 2024. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 13(9): 9.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Liang, J.; He, R.; and Tan, T. 2025. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 133(1): 31–64.
- Liu, Y.; Kothari, P.; Van Delft, B.; Bellot-Gurlet, B.; Mordan, T.; and Alahi, A. 2021. Ttt++: When does self-supervised test-time training fail or thrive? *Advances in Neural Information Processing Systems*, 34: 21808–21820.
- Niu, G.; Dai, B.; Yamada, M.; and Sugiyama, M. 2014. Information-theoretic semi-supervised metric learning via entropy regularization. *Neural computation*, 26(8): 1717–1762.
- Petrov, A.; Torr, P. H.; and Bibi, A. 2023. When do prompting and prefix-tuning work? a theory of capabilities and limitations. *arXiv preprint arXiv:2310.19698*.
- Prabhudesai, M.; Chen, L.; Ippoliti, A.; Fragkiadaki, K.; Liu, H.; and Pathak, D. 2025. Maximizing Confidence Alone Improves Reasoning. *arXiv preprint arXiv:2505.22660*.
- Rein, D.; Hou, B. L.; Stickland, A. C.; Petty, J.; Pang, R. Y.; Dirani, J.; Michael, J.; and Bowman, S. R. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.

Roy, S.; Siarohin, A.; Sangineto, E.; Bulo, S. R.; Sebe, N.; and Ricci, E. 2019. Unsupervised domain adaptation using feature-whitening and consensus loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9471–9480.

Sahoo, P.; Singh, A. K.; Saha, S.; Jain, V.; Mondal, S.; and Chadha, A. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.

Saito, K.; Kim, D.; Sclaroff, S.; Darrell, T.; and Saenko, K. 2019. Semi-supervised domain adaptation via minimax entropy. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8050–8058.

Sun, Y.; Wang, X.; Liu, Z.; Miller, J.; Efros, A.; and Hardt, M. 2020. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, 9229–9248. PMLR.

Wang, D.; Shelhamer, E.; Liu, S.; Olshausen, B.; and Darrell, T. 2020. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*.

Wang, S.; Zhu, Y.; Liu, H.; Zheng, Z.; Chen, C.; and Li, J. 2024. Knowledge editing for large language models: A survey. *ACM Computing Surveys*, 57(3): 1–37.

Wang, Y.; Yang, Q.; Zeng, Z.; Ren, L.; Liu, L.; Peng, B.; Cheng, H.; He, X.; Wang, K.; Gao, J.; et al. 2025. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*.

Xu, F.; Hao, Q.; Zong, Z.; Wang, J.; Zhang, Y.; Wang, J.; Lan, X.; Gong, J.; Ouyang, T.; Meng, F.; et al. 2025. Towards large reasoning models: A survey of reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*.

Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Yang, A.; Zhang, B.; Hui, B.; Gao, B.; Yu, B.; Li, C.; Liu, D.; Tu, J.; Zhou, J.; Lin, J.; et al. 2024. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.

Zhang, Y.; Zhang, Z.; Guan, H.; Cheng, Y.; Duan, Y.; Wang, C.; Wang, Y.; Zheng, S.; and He, J. 2025. No Free Lunch: Rethinking Internal Feedback for LLM Reasoning. *arXiv preprint arXiv:2506.17219*.

Zhao, H.; Liu, Y.; Alahi, A.; and Lin, T. 2023. On pitfalls of test-time adaptation. *arXiv preprint arXiv:2306.03536*.

Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; Dey, A. K.; et al. 2008. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, 1433–1438. Chicago, IL, USA.

Zuo, Y.; Zhang, K.; Sheng, L.; Qu, S.; Cui, G.; Zhu, X.; Li, H.; Zhang, Y.; Long, X.; Hua, E.; et al. 2025. Ttrl: Test-time reinforcement learning. *arXiv preprint arXiv:2504.16084*.