← Go to **AISTATS 2026 Conference** homepage (/group?id=aistats.org/AISTATS/2026/Conference)

3 Versions

CTRLS: Chain-of-Thought Reasoning via Latent State-Transition

(/pdf?id=afld4XGbQe)

Junda Wu (/profile?id=~Junda_Wu1), Yuxin Xiong (/profile?id=~Yuxin_Xiong1), Xintong Li (/profile?id=~Xintong_Li2), Sheldon Yu (/profile?id=~Sheldon_Yu1), Zhengmian Hu (/profile?id=~Zhengmian_Hu1), Tong Yu (/profile?id=~Tong_Yu3), Rui Wang (/profile?id=~Rui_Wang25), Xiang Chen (/profile?id=~Xiang_Chen9), Jingbo Shang (/profile?id=~Jingbo_Shang2), Julian McAuley (/profile?id=~Julian_McAuley1)

25 Sept 2025 (modified: 31 Oct 2025) AISTATS 2026 Conference Submission
Conference, Senior Area Chairs, Area Chairs, Reviewers, Authors
Revisions (/revisions?id=afld4XGbQe)

Abstract:
Chain-of-thought (CoT) reasoning enables large language models (LLMs) to break down complex problems into interpretable intermediate steps, significantly enhancing model transparency and performance in reasoning tasks. However, conventional CoT methods rely on heuristic sampling without structured modeling of reasoning transitions, constraining their ability to systematically explore and discover diverse and effective reasoning trajectories. In this work, we introduce CTRLS, a framework that formulates CoT reasoning as a Markov decision process (MDP) with latent state transitions, enabling principled and state-aware exploration via distributional reinforcement learning. By modelling reasoning actions as explicit probability distributions in latent space, our approach explicitly models epistemic uncertainty, facilitating robust exploration of the reasoning space. As part of our framework, we introduce an on-policy reinforcement learning strategy incorporating epsilon-greedy exploration and entropy-based regularization to iteratively refine latent state transitions without requiring additional fine-tuning of the underlying LLM. Theoretical analyses provide evidence lower bounds (ELBO), theoretically grounding our transition-aware modeling of latent reasoning dynamics. Further experiments demonstrate improvements in reasoning accuracy, diversity, and exploration efficiency across benchmark reasoning tasks.

Submission Policy Confirmation: I have read and adhere to the AISTATS 2026 submission policies on behalf of myself and my co-authors.
Matching System Confirmation: I agree on behalf of myself and my co-authors.
Student Submission: Yes
Reviewer Nomination: Junda Wu (/profile?id=~Junda_Wu1)
Resubmission Confirmation: Yes
Submission Number: 1219

Filter by reply type... | Filter by author... | Search keywords... | Sort: Newest First

Everyone | Program Chairs | Submission1219... | Submission1219 Area...
Submission1219 Authors | Submission1219... | Submission1219... | Submission1219...
Submission1219... | ✖

11 / 13 replies shown

Add: Withdrawal Official Comment

Official Review of Submission1219 by Reviewer 9ygZ

Official Review by Reviewer 9ygZ 11 Nov 2025, 03:52 (modified: 12 Dec 2025, 14:36)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted, Reviewer 9ygZ

Revisions (/revisions?id=wlu6MYOSRJ)

Summary And Contributions:
The manuscript focuses on Chain-of-Thought (CoT) reasoning and proposes CTRLS, a unified, transition-aware framework for modeling CoT capabilities. It recasts reasoning as a latent-state Markov decision process and integrates a stochastic encoder with a state-conditioned U-Net generator under a single, ELBO-driven variational objective that aligns latent transitions with stepwise generation. An on-policy distributional reinforcement learning phase further refines the transition model, enabling robust exploration and controllable reasoning. Overall, the framework offers a principled treatment of reasoning dynamics with potential gains in accuracy, diversity, and interpretability.

Paper Keywords: LLM,COT, RL
Expertise Keywords: LLM reasoning , transformer, efficient Fine-Tuning
Soundness: Correct / minor errors (e.g., typo or errors that do not affect the main results)
Soundness Justification:
The proposed method appears theoretically sound, and the experiments provide credible empirical evidence. However, a notable limitation is that the evaluation is restricted to the selected baselines; this is insufficient. To strengthen the case, the paper should also include comparisons with reinforcement learning approaches—e.g., PPO or related policy-optimization methods—under the same experimental settings.

Significance: Somewhat significant (e.g., theoretical contribution of limited novelty or empirical gains missing baselines/statistical rigor).
Significance Justification:
The paper clearly identifies a salient weakness of conventional CoT methods and introduces a theoretically grounded, empirically validated approach that successfully addresses this limitation. Overall, the proposed solution is both principled and effective in remedying the identified issue.

Novelty: Incremental compared to existing results
Novelty Justification:
While the paper targets a real deficiency in conventional CoT, many works have already sought to remedy it—e.g., Tree-of-Thought–style search and RL-based approaches such as PPO and GRPO, which are known to improve reasoning performance. The manuscript lacks head-to-head comparisons with these families under matched settings, making it difficult to assess the incremental contribution. I consider this omission the paper's most significant weakness and recommend adding strong baselines (ToT/GoT, PPO/GRPO) with comparable budgets and protocols.

Non-conventional Contributions:
N/A

Clarity:
The writing is clear, concise, and easy to follow throughout.

Relation To Prior Work: Some minor related works are missing or described incorrectly
Reproducibility: Sufficient amount of details available for reproducing the main results
Rating: 4: Weak Accept / Weak Reject (Paper has some merits but also notable weaknesses. Reasons to accept and reject are roughly balanced, and the paper may require further evaluation or discussion. Please use sparingly.)
Confidence: 4: Confident, but not absolutely certain. (It is unlikely, but not impossible, that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work.)

Add: Official Comment

Rebuttal by Authors

Rebuttal

by Authors (Junda Wu (/profile?id=~Junda_Wu1), Tong Yu (/profile?id=~Tong_Yu3), Xiang Chen (/profile?id=~Xiang_Chen9), Sheldon Yu (/profile?id=~Sheldon_Yu1), +6 more (/group/edit?id=aistats.org/AISTATS/2026/Conference/Submission1219/Authors))

28 Nov 2025, 13:59 (modified: 01 Dec 2025, 05:17)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Revisions (/revisions?id=jBcuwNG1lN)

Rebuttal:

Soundness

Thank you for the comments regarding Soundness. We would like to clarify that the paper does not aim to introduce a new reinforcement learning algorithm, nor does it position CTRLS as competing with full parameter RL methods such as PPO or GRPO. Throughout the title, abstract, and the list of contributions, our work is framed as proposing a new latent state Markov modelling of chain of thought reasoning that is supported by a principled ELBO based formulation.

The reinforcement learning component in CTRLS is used only as a lightweight mechanism to support robust and explainable exploration of latent transitions. It does not optimize token level policies and it is not intended as an accuracy improving RL method. Since token level PPO and GRPO target a completely different optimization regime, they are not direct baselines for our modelling oriented contribution and operate at a fundamentally different level.

The value of our approach extends beyond empirical performance. It provides a principled, interpretable and mathematically grounded framework for analyzing how reasoning unfolds over time. It also enables examination of transition stability, controllability, and error propagation, which cannot be captured by token level RL or symbolic search alone. This structure offers a foundational tool for systematic study of reasoning behavior in large language models.

We will clarify this positioning more explicitly in the updated version.

Significance

We appreciate the reviewer’s assessment of Significance. CTRLS is not proposed solely as a technique to improve reasoning accuracy. Its main goal is to introduce a new modelling perspective in which chain of thought is represented as transition aware latent dynamics. This perspective is complementary to search based approaches such as ToT and GoT and to RL based approaches such as PPO and GRPO, all of which operate at the symbolic or token policy level.

Our contributions include a latent state abstraction of reasoning steps, an explicit transition model for $P(s_{t+1}|s_t)$, and a unified ELBO objective connecting generative modelling with latent transitions.

Novelty

Thank you for the comments on Novelty. While methods such as ToT, GoT, PPO, and GRPO also aim to improve reasoning, they do so by symbolic branching or token level policy optimization. CTRLS introduces a totally new direction that models chain of thought as a latent state Markov decision process with spectral abstraction, transition aware latent dynamics, and a variational ELBO formulation.

This contribution provides a structural and theoretical lens for understanding reasoning rather than another search or RL optimization technique. Beyond methodological differences, the novelty of CTRLS also lies in offering a modelling framework that makes the internal dynamics of chain of thought explicitly analyzable. By grounding reasoning trajectories in latent transitions with a probabilistic formulation, CTRLS enables study of structure, stability, and failure modes that cannot be accessed through search based or token level RL methods. This conceptual contribution broadens the space of tools available for understanding how reasoning emerges.

To help readers better recognize this, we will add a concise comparison table and expand the related work section to more explicitly position CTRLS relative to search based and RL based families.

Add: **Official Comment**

Replying to Rebuttal by Authors

Thanks for your Rebttal

Official Comment by Reviewer 9ygZ 03 Dec 2025, 13:43

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Comment:
May I ask how your method compares to ToT in terms of computational cost when achieving the same accuracy level? If your approach requires less computation, could you additionally provide an intuitive explanation for why this is the case?

Add: **Official Comment**

Replying to Thanks for your Rebttal

Thank you for the follow up questions

Edit

Official Comment

by Authors (Junda Wu (/profile?id=~Junda_Wu1), Tong Yu (/profile?id=~Tong_Yu3), Xiang Chen (/profile?id=~Xiang_Chen9), Sheldon Yu (/profile?id=~Sheldon_Yu1), +6 more (/group/edit?id=aistats.org/AISTATS/2026/Conference/Submission1219/Authors))

04 Dec 2025, 23:43

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Comment:
We thank the reviewer for the follow up question. Tree of Thought is an inference time prompting strategy that expands CoT into a tree and evaluates multiple branches at each step. In contrast, CTRLS generates a standard linear CoT and updates only a lightweight latent transition policy and a state aware adapter. As a result, its token level computational cost is close to that of vanilla CoT, whereas ToT typically incurs substantially higher cost whenever more than one branch is explored. We also note that CTRLS represents an early foundational effort to model CoT through explicit latent transitions. To establish the modeling framework clearly, this paper focuses on the fundamental single trajectory CoT setting. Our focus in this work is on modeling and learning latent state transitions rather than on search-based prompting, so ToT is complementary rather than a direct baseline. The framework is also compatible with ToT style search: latent transitions can help initialize promising branches or prune unstable ones, we also can extend out assumption on 5.1 to branching latent transition family $P_\theta(s_{i+1} \mid s_i, \mathrm{branch})$. We appreciate the reviewer's question and will refine the relevant discussions in the updated version for clarity.

Add: **Official Comment**

Official Comment by Authors

Edit

Official Comment

by Authors (Junda Wu (/profile?id=~Junda_Wu1), Tong Yu (/profile?id=~Tong_Yu3), Xiang Chen (/profile?id=~Xiang_Chen9), Sheldon Yu (/profile?id=~Sheldon_Yu1), +6 more (/group/edit?id=aistats.org/AISTATS/2026/Conference/Submission1219/Authors))

08 Dec 2025, 12:02

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Comment:
We hope our previous reply clarify your question about Tree of Thought. Please let us know if there's anything else you'd like us to clarify. We really appreciate your feedback.

Add: **Official Comment**

Replying to Official Comment by Authors

Thanks for your reply

Official Comment by Reviewer 9ygZ 08 Dec 2025, 17:55

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Comment:
I now have a better understanding of your method and am willing to adjust my score upward accordingly.

Add: **Official Comment**

Official Review of Submission1219 by Reviewer HmPF

Official Review by Reviewer HmPF 10 Nov 2025, 01:07 (modified: 12 Dec 2025, 14:36)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted, Reviewer HmPF

Revisions (/revisions?id=ZTPEvdyHVd)

Summary And Contributions:
The paper converted chain-of-thought (CoT) reasoning as a Markov decision process (MDP) with latent state transitions, and then proposed a Chain-of-Thought Reasoning via Latent State Transition framework, i.e., the CTRLS. The presented CTRLS is a unified variational model that jointly learns a latent encoder, a state transition model, and an LLM reasoning adapter, optimized under a single evidence lower bound (ELBO) objective.

The key contributions of this paper include:

1. A novel latent MDP formulation for explicit modeling of complex state transitions among multiple reasoning steps
2. A distributional reinforcement learning strategy for exploring CoT inference
3. Theoretical derivation of finite-sample evidence lower bounds (ELBO) for training
4. Empirical evaluations of the improved CoT reasoning in different perspectives on math benchmarks.

Paper Keywords: Chain-of-thought reasoning, Markov decision process, distributional reinforcement learning,
Expertise Keywords: Large language models, Chain-of-thought reasoning, supervised learning
Soundness: Correct / minor errors (e.g., typo or errors that do not affect the main results)
Soundness Justification:
The theoretical framework is basically sound, including the MDP formulation and proper ELBO derivations. However, there are several concerns:

1. The binary reward function (1{y = y*}) in the on-policy RL is too simple for complex reasoning states and may discourage some correct reasoning states followed by very few incorrect ones.
2. The spectrum clustering for latent state extraction lacks enough support, i.e., theoretical proofs or empirical validation, for why this particular representation would capture reasoning semantics accurately.
3. The Markov assumption for reasoning transitions may be too strong for complex reasoning tasks, as (long) Chain-of-Thought reasoning tends to involve long-distance dependencies.

Significance: Somewhat significant (e.g., theoretical contribution of limited novelty or empirical gains missing baselines/statistical rigor).
Significance Justification:
This work interpreted the current CoT reasoning mechanism by introducing structured transition modeling, and the integration of latent states with distributional RL for reasoning is quite novel and interesting. The empirical improvements on 2 math benchmarks demonstrated its effectiveness, which illustrated consistent gains across different LLMs and datasets.

Novelty: New results
Novelty Justification:

While MDP formulations have been employed in RL for LLMs in previous works, applying distributional RL with latent state transitions to CoT reasoning is still novel. For example, The Coconut [1] explored continuous latent reasoning but didn't provide the structured transition modeling and explicit exploration mechanisms like the presented CTRLS.

[1] Training Large Language Models to Reason in a Continuous Latent Space. COLM 2025.

Non-conventional Contributions:
The paper bridged the reinforcement learning theory with LLM CoT reasoning, stepping beyond simple prompting or vanilla fine-tuning approaches. The employment of distributional policies for reasoning actions is unconventional in the LLM reasoning research. The two-phase training strategy that combined offline pretraining with on-policy RL is also novel for math reasoning tasks.

Clarity:
The paper is generally well-written with clear mathematical notations and strict theoretical derivations. However, some sections could still be improved:

1. The relation between the latent states and concrete semantics of reasoning states needs better explanation.
2. The ablation studies could be more systematic in isolating the contributions of different components, and the authors could consider broader scales of LLMs and reasoning tasks for evaluating the proposed CTRLS.
3. The computational costs and scalability considerations were not discussed exhaustively.

Relation To Prior Work: All related works are clearly discussed
Reproducibility: Some amount of details available
Rating: 6: Accept (Technically solid paper, with high impact on at least one sub-area. The contribution is convincing and the evaluation is adequate, though there may be some limitations or open questions.)
Confidence: 4: Confident, but not absolutely certain. (It is unlikely, but not impossible, that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work.)

Add: **Official Comment**

Rebuttal by Authors

Rebuttal

by Authors (Junda Wu (/profile?id=~Junda_Wu1), Tong Yu (/profile?id=~Tong_Yu3), Xiang Chen (/profile?id=~Xiang_Chen9), Sheldon Yu (/profile?id=~Sheldon_Yu1), +6 more (/group/edit?id=aistats.org/AISTATS/2026/Conference/Submission1219/Authors))

28 Nov 2025, 14:03 (modified: 01 Dec 2025, 05:17)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Revisions (/revisions?id=MRSi1uTxwt)

Rebuttal:
We thank the reviewer for the positive and encouraging evaluation, and we appreciate the thoughtful feedback. Below we address the specific concerns raised.

1. Binary reward and partially correct trajectories

We appreciate the reviewer's concern. In CTRLS, the binary reward is not used to shape token-level behavior but only to update the latent transition matrix. Since transitions operate at the latent–state level, a coarse terminal reward is sufficient and aligns with common practice in latent-policy RL for reasoning (e.g., GRPO-style and RLAE-style MDP formulations). We will clarify this design choice in the revision.

2. Spectrum clustering for latent state extraction

Our work use spectral clustering is (i)To capture the evolving semantics of each reasoning step and (ii)To incorporate contextual accumulation and recursively update the gram matrix. This emable us to capture the CoT clustering in the gram matrix. We adopt the methodology from spectral clustering (Ng et al.(2001) Macgregor(2023) Laenen et al.(2024)). CTRLS is a framework that abstracts CoT reasoning as structured transitions over latent semantic states, models them based of Markov chain.

[1] On Spectral Clustering: Analysis and an Algorithm, 2001, A.Y. Ng, M.I. Jordan, Y. Weiss

[2] Fast and Simple Spectral Clustering in Theory and Practice, 2023, P. Macgregor

[3] Dynamic Spectral Clustering with Provable Approximation Guarantee, 2024, S. Laenen, H. Sun

3. Markov assumption for long reasoning chains

We appreciate the reviewer's comment. Our Markov assumption is applied only at the *latent* level, where z_t summarizes the entire prefix $x_{\leq t}$, so long-range dependencies are preserved. This abstraction is standard in RL-for-LLM settings used by GRPO and PPO-based RLHF, and many subsequent works—such as RLAE [1]—adopt the same MDP formulation.

[1] RLAE: Reinforcement Learning-Assisted Ensemble for LLMs. EMNLP 2025.

4. Clarifications on presentation

We agree that more explicitly organised ablations would make the contributions clearer. Several components are already ablated in the current experiments: (i) Table 1 compares CTRLS against the backbone LLMs under identical sampling budgets, effectively ablating the transition-aware module; (ii) Table 2 together with Fig. 4–5 systematically varies the entropy weight ($H \in \{0, 0.01\}$) and the ϵ-greedy rate ($\epsilon \in \{0, 0.1, 0.3\}$), isolating the effects of entropy regularisation and stochastic exploration in our on-policy RL phase; and (iii) Table 3 evaluates the same configurations under a GPT-4 rubric (self-reflection, algebra, coherence, hallucinations), showing consistent trends.

Thank you again for the constructive comments and helpful suggestions.

Add: **Official Comment**

Replying to Rebuttal by Authors

Official Comment by Reviewer HmPF

Official Comment by Reviewer HmPF 08 Dec 2025, 18:27

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Comment:
Thanks for the author's exhaustive responses! I have read all the responses, and most of my concerns were addressed, so I decided to raise my rating from 5 to 6.

However, I still recommend supplementing experiments with more models of different scales (e.g., 7B/8B/14B) and more benchmarks (e.g., AIME/GPQA-Diamond/AMC) to demonstrate the scalability.

Add: **Official Comment**

Official Review of Submission1219 by Reviewer Z5sQ

Official Review by Reviewer Z5sQ 15 Oct 2025, 14:36 (modified: 12 Dec 2025, 14:36)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted, Reviewer Z5sQ

Revisions (/revisions?id=JrGKP7mf5i)

Summary And Contributions:
The paper proposes CTRLS, a framework that formulates chain-of-thought (CoT) reasoning as latent state transitions within a Markov decision process (MDP). A variational encoder maps partial CoT prefixes to latent states; a transition model outputs distributional actions (Dirichlet) over next-state proposals; and an adapted LLM emission model conditions token generation on these latent states. Training proceeds in two phases: (i) offline pre-training via an ELBO that jointly optimizes the encoder, transition model, and emission adapter; and (ii) on-policy reinforcement learning with entropy regularization and ε-greedy exploration, optimizing only the transition policy for task accuracy. Experiments on GSM8K and MATH with LLaMA-3.2-3B-Instruct and Qwen2.5-3B-Instruct report gains in exploration accuracy (pass@20) and success rate across temperatures/ε settings, plus qualitative improvements in self-reflection and algebraic correctness.

Paper Keywords: CoT, LLMs, RL

Expertise Keywords: XAI, LLMs, VLMs
Soundness: Correct / minor errors (e.g., typo or errors that do not affect the main results)
Soundness Justification:
Mostly sound, with some open assumptions and a few places needing clarification.

- The **variational formulation and ELBO** are standard and correctly specialized to their factorization; the per-time-step decomposition and KL terms are consistent.
- The **MDP abstraction** plus **Dirichlet policy** is coherent for exploration in a learned latent CoT space. The epsilon-greedy mixing and entropy bonus are routine and appropriate.
- The **state-aware adapter** integrates z_t into token features via a bottleneck U-Net; the dependency $P_\omega(x_t \mid x_{<t}, z_{\leq t})$ vs. the Markov simplification $P_\omega(x_t \mid x_{<t}, z_{t-1})$ is plausible but would benefit from tighter notation.

Significance: Significant contributions (e.g., strong theoretical insights or well-supported empirical improvements with appropriate baselines/statistics).
Significance Justification:
CTRLS pushes CoT beyond token-level heuristics, offering a principled control handle (latent transitions) that can be optimized with RL without touching the base model weights. If robust, this could be a useful ingredient for controllable reasoning and for integrating exploration with state-aware adapters. The conceptual shift—transition modeling in latent CoT space—is meaningful and could generalize to other reasoning domains (code, science QA). The empirical gains are consistent but modest and measured on limited settings; broader benchmarks and stronger baselines would lift impact.

Novelty: Ground-breaking
Novelty Justification:
CTRLS differs from prior CoT sampling/search by: (i) latent-state MDP for reasoning; (ii) distributional (Dirichlet) policies over latent actions; (iii) a joint ELBO to align encoder/transition/generator before on-policy RL. While latent reasoning and continuous-space ideas exist, the particular combo with transition policies and frozen-backbone, state-aware adapters appears new.

Non-conventional Contributions:
Bridging CoT with distributional RL in a learned latent space and using a light U-Net adapter to inject state into token features are unconventional design choices for reasoning with frozen LLMs.

Clarity:
Mostly clear

- Contributions and high-level story are clearly stated;
- Mathematical sections define assumptions and give an explicit ELBO/Theorem.

Relation To Prior Work: All related works are clearly discussed
Reproducibility: Some amount of details available
Rating: 6: Accept (Technically solid paper, with high impact on at least one sub-area. The contribution is convincing and the evaluation is adequate, though there may be some limitations or open questions.)
Confidence: 3: Fairly confident (It is possible that you did not understand some parts of the submission or that you are unfamiliar with some pieces of related work. Math/other details were not carefully checked.)

Add: **Official Comment**

Rebuttal by Authors

Rebuttal

by Authors (Junda Wu (/profile?id=~Junda_Wu1), Tong Yu (/profile?id=~Tong_Yu3), Xiang Chen (/profile?id=~Xiang_Chen9), Sheldon Yu (/profile?id=~Sheldon_Yu1), +6 more (/group/edit?id=aistats.org/AISTATS/2026/Conference/Submission1219/Authors))

28 Nov 2025, 14:00 (modified: 01 Dec 2025, 05:17)

Program Chairs, Authors, Senior Area Chairs, Area Chairs, Reviewers Submitted

Revisions (/revisions?id=wlVtQb4mqG)

Rebuttal:
We thank the reviewer for the positive and thoughtful evaluation.

Soundness. We appreciate the assessment that the ELBO, MDP abstraction, and factorization choices are sound. We also thank the reviewer for pointing out that the notation surrounding the state-aware adapter and its Markov simplification could be tightened. We will revise the notation and provide a clearer description in the camera-ready to improve readability and precision.

Significance. We are grateful that the reviewer recognized the value of framing CoT reasoning through latent-state transitions and noted its broader applicability beyond math domains.

Novelty. We appreciate the recognition that combining latent-state MDPs, distributional transition policies, and a jointly trained encoder/transition/generator is a distinct and unconventional design.

Thank you again for the encouraging feedback.

Add: **Official Comment**