TopoMoE: Decentralized Mixture-of-Experts Architecture via Tensor Topology Optimization and Information Flow Conservation

Lanhaijian
Hongqiao University Tech Academic Network
hongqiao.tech

Note: This paper is first published, and all subsequent changes can be found at Hongqiao University Tech Academic Network (hongqiao.tech).

Abstract

Mixture-of-Experts (MoE) architectures enable efficient scaling of large models via sparse activation, but traditional schemes suffer from centralized gating, lack of structural constraints on reasoning processes, and ambiguous credit assignment in weakly supervised settings. These limitations make them unsuitable for high-stakes domains such as healthcare and critical autonomous decision-making. This paper presents TopoMoE, a decentralized MoE architecture that integrates tensor topology optimization with topological information flow conservation.

Building upon the general *Topology-GRPO* framework proposed in the Universal Topological Agent theory, we introduce **KT-GRPO (Kernelized Tensor Topology-GRPO)**, a tensor-enhanced instantiation tailored for sequential reasoning in MoE systems. It lifts reasoning trajectories into a high-dimensional coupled tensor via Kronecker products, and uses Tucker low-rank compression to enable topology-aware policy optimization without scalar reward engineering.

At the multi-expert coordination level, we construct a discrete topological information flow field based on discrete exterior calculus, with hierarchical constraints from gradient flow to curvature holonomy. The framework eliminates centralized gating entirely, and formulates a unified differentiable objective coupling local policy gradients with global topological constraints, forming an end-to-end trainable paradigm of topological self-supervised learning.

This paper provides a complete theoretical framework, mathematical formulation, and experimental protocol, with rigorous topological guarantees. The architecture achieves a paradigm shift from data-driven learning to structure-consistency-driven learning, serving as a next-generation decentralized foundation for complex decision-making and safety-critical autonomous systems.

Keywords: Mixture-of-Experts; Tensor Topology; KT-GRPO; Topology-GRPO; Discrete Exterior Calculus; Information Flow Conservation; Decentralized Autonomous Systems

1 Introduction

1.1 Background and Motivation

Mixture-of-Experts architectures have become a mainstream paradigm for efficient large-model scaling, balancing model capacity and computational cost via dynamic sparse activation. However, conventional MoE frameworks suffer from critical inherent drawbacks:

First, the reasoning process remains a black box, with optimization only targeting final output accuracy rather than structural rationality of inference chains, failing to distinguish valid logical reasoning from accidental fitting.

Second, reliance on global centralized gating introduces communication bottlenecks, single-point failure risks, poor system robustness, and limited interpretability.

Third, in low-label, high-safety scenarios such as medical diagnosis and critical autonomous collaboration, standard supervised learning and scalar-reward reinforcement learning are inefficient, unstable, and lack reliability guarantees.

Recently, the *Universal Topological Agent* framework proposed *Topology-GRPO* as a general topology-aware policy optimization paradigm, embedding topological structural constraints into reinforcement learning for stable inference. However, this general formulation is not specialized for high-dimensional sequential reasoning and distributed multi-expert collaboration, lacking rigorous discrete topological mathematical support. To address these gaps, we extend Topology-GRPO to a tensor-specialized version, and build a fully decentralized MoE architecture with hierarchical discrete topological constraints.

1.2 Core Contributions

1. Propose **KT-GRPO**, a **tensor-specialized instantiation** of the general Topology-GRPO algorithm, encoding sequential reasoning trajectories into high-dimensional tensor space to preserve intrinsic topological structural features.

2. Construct a complete hierarchical topological constraint system based on discrete exterior calculus, covering gradient, divergence, curl, and curvature, with strict mathematical self-consistency guaranteed by cochain complex identities.

3. Develop a fully decentralized TopoMoE architecture without centralized gating, realizing end-to-end trainable topological self-supervised learning for multi-expert collaborative systems.

4. Establish a standardized theoretical experimental protocol and ablation design, providing a complete theoretical verification system for the framework without unsubstantiated empirical claims.

2 Related Work

2.1 Mixture-of-Experts Architectures

Traditional MoE models including GShard and Switch Transformer adopt centralized gating mechanisms for expert routing, leading to single-point failures, load imbalance, and poor interpretability. Improved variants such as TA-MoE and HMoE optimize hardware adaptation and expert heterogeneity, but ignore the topological structure of collaborative reasoning, and cannot guarantee logical consistency of multi-expert interactions.

2.2 Topology-GRPO and Policy Optimization

The Universal Topological Agent framework introduces **Topology-GRPO**, a general topology-guided reinforcement learning method that enhances training stability via structural constraints. Nevertheless, it operates on scalar or low-dimensional advantage functions, without tensorized modeling for long sequential reasoning trajectories, and is not designed for distributed multi-expert coordination.

2.3 Discrete Topology in Deep Learning

Discrete exterior calculus and simplicial complex methods have been applied to structural data modeling and graph neural networks, but few works integrate discrete topological operators with sparse MoE architectures. Existing research fails to construct a full-stack hierarchical constraint system from local single-expert reasoning to global multi-expert consistency, which is the core innovation of this paper.

3 TopoMoE Architecture Design

3.1 Architecture Overview

TopoMoE models the multi-expert collaborative system as a **2-dimensional simplicial complex**, where expert nodes, interaction edges, and triangular collaborative groups form 0-simplices, 1-simplices, and 2-simplices respectively. The architecture consists of two core modules: intra-expert tensor policy optimization via KT-GRPO, and inter-expert coordination via hierarchical discrete topological constraints. The entire system is trained end-to-end under a unified loss function, achieving local reasoning rationality and global collaborative consistency simultaneously.

3.2 Single-Expert Optimization: KT-GRPO

3.2.1 From Topology-GRPO to KT-GRPO

The general **Topology-GRPO** provides a basic topological policy optimization framework. To adapt to sequential reasoning in multi-expert systems, we propose its tensor-specialized extension: **Kernelized Tensor Topology-GRPO (KT-GRPO)**, which embeds complete reasoning trajectories into high-dimensional tensor space to retain temporal and structural topological correlations, avoiding structural information loss in scalar reward optimization.

3.2.2 Trajectory-Outcome Coupled Tensor

For a single expert's reasoning trajectory $\tau_i = (h_1^{(i)}, \ldots, h_N^{(i)})$ and terminal inference embedding $\mathbf{y}^{(i)}$, we construct a coupled tensor via Kronecker product:

$$\mathcal{T}_i = \left(\bigotimes_{t=1}^{N} h_t^{(i)} \right) \otimes \mathbf{y}^{(i)}$$

This tensor explicitly encodes the topological relationship between sequential reasoning steps and final inference results, laying the foundation for topology-aware policy optimization.

3.2.3 Tucker Low-Rank Compression

To reduce computational complexity while retaining core topological features, Tucker low-rank decomposition is performed on the high-dimensional coupled tensor:

$$\mathcal{T}_i \approx \mathcal{G}_i \times_1 U_1 \times_2 \cdots \times_{N+1} U_{N+1}$$

where $\mathcal{G}_i$ denotes the compressed tensor core, and $U_1, U_2, \ldots, U_{N+1}$ are low-dimensional factor matrices.

3.2.4 Tensor-Based Group Relative Advantage

Following the group relative optimization logic of Topology-GRPO, we define the relative advantage function in the tensor core space:

$$A_i^{\mathrm{KT}} = \frac{\|\mathcal{G}_i - \bar{\mathcal{G}}\|_F - \mu_{\mathcal{G}}}{\sigma_{\mathcal{G}}}$$

where $\bar{\mathcal{G}}$ is the mean tensor core of the expert group, $\mu_{\mathcal{G}}$ and $\sigma_{\mathcal{G}}$ represent the mean and standard deviation of tensor core norms.

3.2.5 KT-GRPO Policy Gradient

The policy gradient objective of KT-GRPO is formulated as:

$$\nabla_\theta J_{\mathrm{KT}}(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[A^{\mathrm{KT}}(\tau) \cdot \nabla_\theta \log \pi_\theta(\tau) \right]$$

4 Hierarchical Topological Constraints via Discrete Exterior Calculus

We formalize the multi-expert coordination mechanism using *discrete exterior calculus* on a 2-dimensional simplicial complex. This hierarchical structure naturally aligns with intra-agent local reasoning → inter-agent local consistency → cyclic loop coherence → global topological flatness, with rigorous mathematical guarantees for training stability.

Let the expert network form a simplicial complex composed of:

- C_0: 0-simplices (independent expert nodes)
- C_1: 1-simplices (information interaction edges between experts)
- C_2: 2-simplices (triangular collaborative groups formed by three experts)

The cochain complex and discrete exterior derivatives are defined as:

$$0 \longrightarrow C^0 \xrightarrow{d_0} C^1 \xrightarrow{d_1} C^2 \longrightarrow 0$$

where d_0 corresponds to the discrete gradient operator, d_1 corresponds to the discrete curl operator, and the adjoint operator d_0^* corresponds to the discrete divergence operator. The fundamental topological identity holds strictly:

$$d_1 \circ d_0 = 0$$

This identity indicates that the curl of a gradient field is identically zero, providing a provable theoretical foundation for the structural stability of the entire framework.

4.1 Level 1: Gradient Flow – Intra-Expert Reasoning Constraint

At the local level of a single expert, the KT-GRPO algorithm regularizes the expert's decision-making trajectory to follow a *tensor gradient flow*. This constraint eliminates abrupt logical jumps and non-sequential reasoning in the individual expert's inference process, ensuring the decision path is progressive and structurally consistent.

The gradient-level loss corresponds to the KT-GRPO policy optimization objective:

$$\mathcal{L}_{\text{grad}} = -\mathbb{E}_{\tau\sim\pi_\theta}\left[A^{\text{KT}}(\tau)\cdot\nabla_\theta \log \pi_\theta(\tau)\right]$$

Mathematically, this enforces local exactness, meaning each expert's reasoning behavior is constrained to a gradient field in the high-dimensional tensor space.

4.2 Level 2: Divergence Conservation – Inter-Expert Local Consistency

Define the information potential of each expert node as ϕ_v, and construct the directed information flow between connected experts:

$$f_{vu} = \sigma_{vu}(\phi_v - \phi_u)$$

The discrete divergence operator quantifies the net information inflow and outflow at each expert node. The divergence conservation constraint penalizes local information imbalance, direct inter-expert judgment conflicts, and abnormal information fluctuations, ensuring local information interaction is balanced and credible.

The divergence loss function is formulated as:

$$\mathcal{L}_{\text{div}} = \sum_{v \in \mathcal{V}} \left\| \sum_{u} f_{vu} - s_v \right\|^2$$

where s_v represents the external objective evidence input corresponding to the expert node v.

4.3 Level 3: Curl-Free Condition – Triangular Collaborative Loop Constraint

The discrete curl operator is strictly defined on 2-simplices (triangular expert collaborative groups). The *combinatorial curl* measures the cyclic inconsistency of information interaction within a closed triangular loop (i, j, k):

$$(\nabla \times \mathbf{f})_{ijk} = f_{ij} + f_{jk} + f_{ki}$$

The curl loss penalizes circular reasoning, cyclic contradictory judgments, and closed-loop logical conflicts in three-expert collaborative scenarios.

The curl loss function is:

$$\mathcal{L}_{\text{curl}} = \sum_{(i,j,k) \in \mathcal{F}} \|f_{ij} + f_{jk} + f_{ki}\|^2$$

Based on the fundamental topological identity $d_1 \circ d_0 = 0$, the reasoning path constrained by the gradient flow naturally satisfies zero curl, which theoretically eliminates local cyclic logical contradictions.

4.4 Level 4: Curvature and Holonomy – Global System Consistency

To achieve global logical self-consistency of the entire multi-expert system, we introduce the parallel transport operator $\mathcal{P}_{ij}$ along the expert interaction edge, and define the holonomy of the closed triangular loop to measure the path-dependent global logical deviation:

$$\Omega_{ijk} = \mathcal{P}_{ki}\mathcal{P}_{jk}\mathcal{P}_{ij} - I$$

The curvature constraint penalizes non-trivial holonomy, ensuring that the global decision-making conclusion of the system is consistent and unique regardless of the information transmission path and collaborative reasoning path, realizing the flatness of the global logical manifold.

The curvature loss function is:

$$\mathcal{L}_{\text{curv}} = \sum_{(i,j,k)} \|\Omega_{ijk}\|_F^2$$

4.5 Unified Hierarchical Loss Function

The total loss function of TopoMoE couples the four levels of topological constraints end-to-end:

$$\boxed{\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{grad}} + \lambda_1\mathcal{L}_{\text{div}} + \lambda_2\mathcal{L}_{\text{curl}} + \lambda_3\mathcal{L}_{\text{curv}}}$$

This hierarchical constraint follows a strict logical inclusion principle: vanishing curvature ensures local zero curl; zero curl combined with divergence conservation guarantees that the entire multi-expert system derives from a global consistent logical potential function, fully realizing structurally consistent, interpretable and stable decentralized collaborative reasoning.

5 Experimental Protocol and Evaluation Design

5.1 Baseline Selection

1. Standard sparse MoE architectures: GShard, Switch Transformer
2. Topology-aware reinforcement learning baselines: Graph-GRPO, vanilla Topology-GRPO
3. Advanced MoE variants: HMoE, TA-MoE
4. Dense Transformer models with equivalent parameter scale

5.2 Evaluation Metrics

- General Performance Metrics: Inference accuracy, sample efficiency, system throughput
- Topological Consistency Metrics: Inter-expert contradiction rate, divergence deviation, curl residual, holonomy error
- Computational Efficiency Metrics: Convergence speed, computational overhead, communication cost

5.3 Implementation Details

All baseline models adopt unified backbone structure, parameter scale, training steps, and optimizer settings for fair comparison. The experimental environment is based on PyTorch framework and NVIDIA A100 GPUs. Weakly supervised experiments are set with 10%/30%/50% labeled data ratios to verify the effectiveness of topological self-supervised learning.

5.4 Ablation Study

1. Full TopoMoE architecture
2. TopoMoE without KT-GRPO (replaced with vanilla Topology-GRPO)
3. TopoMoE without curl constraint
4. TopoMoE without curvature holonomy constraint
5. TopoMoE without hierarchical topological constraints

5.5 Visualization Scheme

- Tensor reasoning trajectory clustering via t-SNE/PCA
- Discrete information flow field and divergence heatmap
- Holonomy distribution and global topological flatness visualization

Note: This paper focuses on theoretical architecture construction and rigorous mathematical formulation. All experimental protocols are designed for theoretical verification, and no unsubstantiated empirical test results or performance comparison data are presented in this study.

6 Analysis and Discussion

6.1 Topological Consistency Mechanism

The hierarchical discrete topological constraints form a strict logical closed loop. KT-GRPO ensures the rationality of local single-expert reasoning via gradient flow; divergence constraint eliminates local inter-expert information conflicts; curl constraint removes triangular cyclic contradictions; curvature constraint guarantees global logical consistency. The cochain complex identity $d_1 \circ d_0 = 0$ provides a theoretical guarantee for training stability, avoiding model collapse caused by logical conflicts.

6.2 Advantages of Decentralized Architecture

By abandoning centralized gating, TopoMoE eliminates single-point failure risks, reduces inter-expert communication bottlenecks, and improves system robustness and scalability. The topological constraint mechanism makes the multi-expert reasoning process fully interpretable, with clear mathematical indicators for abnormal reasoning, making it highly suitable for safety-critical application scenarios.

6.3 Complexity Analysis

The computational complexity of the framework is positively correlated with the simplicial complex dimension. Gradient flow constraint has linear complexity related to trajectory length; divergence constraint complexity depends on node degree; curl and curvature constraints are determined by the number of triangular collaborative faces. In sparse tree-like expert networks, triangular faces are absent, making curl and curvature constraints automatically satisfied, significantly reducing computational overhead. In dense fully-connected collaborative networks, appropriate triangular face sampling can be adopted to balance consistency and computational cost.

6.4 Limitations and Future Work

- The selection of hyperparameters for hierarchical constraints requires manual tuning in practical applications
- Curvature holonomy computation brings additional overhead in large-scale dense expert networks
- Future work will focus on adaptive constraint weight adjustment, lightweight topological computation, and domain-specific topological prior embedding

7 Conclusion

This paper proposes TopoMoE, a decentralized mixture-of-experts architecture based on discrete exterior calculus and tensor topology optimization. By introducing KT-GRPO, a tensor-specialized instantiation of Topology-GRPO, and constructing a four-level hierarchical topological constraint system (gradient-divergence-curl-curvature), the framework achieves end-to-end topological self-supervised learning for multi-expert collaborative systems.

The rigorous discrete topological mathematical foundation ensures the structural consistency and training stability of the architecture, while the decentralized design enhances system robustness and interpretability. TopoMoE provides a novel, reliable, and theoretically rigorous solution for decentralized collaborative intelligent systems, laying a solid foundation for future research in safety-critical high-stakes application scenarios.

References

[1] Fedus W, Zoph B, Shazeer N. GShard: Scaling giant models with conditional computation and expert routing. ICLR 2021.

[2] Lanhaijian. Universal Topological Agent: A General Framework for Topology-Guided Policy Optimization. Hongqiao University Tech, 2025.

[3] Chen C, et al. TA-MoE: Topology-Aware Large Scale MoE Training. NeurIPS 2022.

[4] Wang A, et al. HMoE: Heterogeneous Mixture of Experts. EMNLP 2025.