Diversity and Coordination in LLM Reasoning Traces: A Computational Study of Implicit Multi-Perspective Problem Solving

Anonymous Author(s)

ABSTRACT

Modern reasoning-optimized large language models (LLMs) implicitly simulate multi-agent-like dialogue among diverse internal perspectives during chain-of-thought reasoning, yet the mechanisms governing diversity and coordination within these traces remain unresolved. We formalize this open question through a tractable surrogate framework that models reasoning traces as paths in a multi-modal solution landscape with distinct optima of varying quality and difficulty. We compare four coordination mechanisms: Independent sampling, Repulsive Sampling (RS) with RBF diversity kernels, Strategy-Conditioned Generation (SCG), and Ensemble Coordination (EC) with specialized sub-policies. Across 50 evaluation problems with 80 coordination rounds, our experiments reveal that SCG achieves the highest strategy coverage (0.794 vs. 0.387 for Independent), while RS attains the best peak solution quality (max quality 0.555 vs. 0.482 for Independent). We observe a fundamental diversity–accuracy tradeoff: methods maximizing endpoint diversity (cosine diversity 0.918 for SCG) do not always maximize solution quality, suggesting that *structured* diversity—targeting distinct solution strategies rather than maximizing geometric spread—is key to effective coordination. Ablation studies over trace count $K \in \{2, 4, 8, 16\}$ and repulsion strength confirm that coordination benefits scale with K and exhibit a sweet spot for repulsion intensity. These findings provide a quantitative framework for understanding how implicit perspectives within LLM reasoning traces can be organized to improve collective problem solving.

CCS CONCEPTS

• **Computing methodologies → Neural networks**; *Learning latent representations.*

KEYWORDS

diversity, coordination, chain-of-thought reasoning, large language models, multi-agent systems

ACM Reference Format:

Anonymous Author(s). 2026. Diversity and Coordination in LLM Reasoning Traces: A Computational Study of Implicit Multi-Perspective Problem Solving. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 4 pages.

1 INTRODUCTION

Recent work has demonstrated that reasoning-optimized LLMs such as DeepSeek-R1 and QwQ-32B implicitly generate multi-agent-like conversational structures within their chain-of-thought traces [7]. These models exhibit conversation-like behaviors including question-answering exchanges, perspective shifts, and conflict-reconciliation patterns that resemble a "society of thought" operating within a single model's reasoning process. This finding connects to a broader question in collective intelligence: how does diversity among problem-solving perspectives, combined with coordination mechanisms, affect the quality of solutions discovered [5, 10]?

While chain-of-thought prompting [12] and self-consistency decoding [11] have shown that generating multiple reasoning traces and aggregating their outputs improves accuracy, these methods treat traces as independent samples. More structured approaches such as Tree of Thoughts [13] and Graph of Thoughts [1] impose explicit structure on the reasoning process, while multi-agent debate frameworks [4, 9] assign distinct roles to separate model instances. However, how diversity and coordination operate *within* the internal reasoning traces of a single LLM remains an open question [7].

We address this question through a computational framework that models LLM reasoning traces as paths in a multi-modal solution landscape. Our contributions are:

1. A **formal surrogate framework** modeling reasoning traces as paths through a combinatorial space with multiple solution optima of varying quality and difficulty.
2. A **systematic comparison** of four coordination mechanisms—Independent, Repulsive Sampling, Strategy-Conditioned Generation, and Ensemble Coordination—measuring their effects on diversity and solution quality.
3. **Quantitative evidence** that structured diversity (strategy coverage) matters more than geometric diversity (cosine distance) for solution quality, revealing a nuanced diversity–accuracy tradeoff.
4. **Ablation studies** characterizing how trace count and repulsion strength modulate coordination effectiveness.

2 RELATED WORK

Chain-of-Thought Reasoning. Wei et al. [12] demonstrated that prompting LLMs to produce intermediate reasoning steps substantially improves performance on complex tasks. Self-consistency [11] extended this by sampling multiple reasoning paths and selecting the most common answer, implicitly leveraging diversity. Step-aware verification [8] further refines trace quality through process-level supervision.

Structured Reasoning. Tree of Thoughts [13] and Graph of Thoughts [1] impose explicit graph structures on reasoning, enabling backtracking and combination of partial solutions. These approaches demonstrate that structured exploration of the reasoning space yields benefits beyond independent sampling.

Multi-Agent LLM Systems. Multi-agent debate [4, 9] assigns distinct personas to multiple LLM instances, and frameworks like MetaGPT [6] formalize role differentiation. The "society of thought" perspective [7] suggests that similar dynamics emerge implicitly within single-model reasoning traces.

Mechanistic Interpretability of Reasoning. Sparse autoencoders [2, 3] have enabled identification of interpretable features within LLM activations, including features associated with reasoning patterns. Kim et al. [7] used such tools to demonstrate that diversity-related features are causally linked to reasoning performance.

3 METHODS

3.1 Problem Landscape Model

We model each reasoning problem as a multi-modal optimization landscape in $\mathbb{R}^D$ (with $D = 20$) containing $M = 6$ solution optima. Each optimum j is characterized by a center $\mathbf{c}_j$, a base quality $q_j = 0.3 + 0.12j$, a basin width $w_j = \max(2.0 - 0.2j, 0.5)$, and a difficulty parameter $\delta_j = 0.2 + 0.15j$. Higher-indexed optima represent harder but more rewarding solution strategies.

A reasoning trace $\tau = (\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_T)$ is a path of length $T = 12$ through this space. The quality of a trace is evaluated at its endpoint:

$$Q(\tau) = \max_{j \in [M]} q_j \cdot \exp\left(- \frac{\|\mathbf{x}_T - \mathbf{c}_j\|^2}{2w_j^2} \right) \cdot \frac{1}{1 + \delta_j \|\mathbf{x}_T - \mathbf{c}_j\|} \quad (1)$$

plus Gaussian noise $\mathcal{N}(0, 0.15^2)$. This creates a diversity–accuracy tradeoff: easy optima (low j) have broad basins but low quality, while hard optima (high j) have narrow basins and high quality.

3.2 Coordination Mechanisms

Given $K = 8$ traces per problem, we compare four mechanisms:

Independent Sampling (Baseline). Each trace follows a biased random walk toward the nearest optimum with step $\mathbf{s}_t = 0.3\hat{\mathbf{d}}_{\text{nearest}} + 0.5\boldsymbol{\epsilon}_t$, where $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Repulsive Sampling (RS).. A diversity kernel repels each trace from previously generated traces using an RBF force:

$$\mathbf{f}_{\text{repel}}(\mathbf{x}_t) = \lambda \sum_{i<k} \exp\left(- \frac{\|\mathbf{x}_t - \tau_i(t)\|^2}{2h^2} \right) \cdot \frac{\mathbf{x}_t - \tau_i(t)}{\|\mathbf{x}_t - \tau_i(t)\|} \quad (2)$$

with repulsion strength $\lambda = 0.5$ and bandwidth $h = 1.5$.

Strategy-Conditioned Generation (SCG).. Each trace k is assigned strategy label $s_k = k \bmod M$ and biased toward the corresponding optimum with enhanced attraction: $\mathbf{s}_t = 0.6\hat{\mathbf{d}}_{s_k} + 0.15\hat{\mathbf{d}}_{\text{nearest}} + 0.35\boldsymbol{\epsilon}_t$.

Ensemble Coordination (EC).. A portfolio of $E = 4$ sub-policies, each with a learned directional bias $\mathbf{b}_e$, assigns traces round-robin and updates biases based on strategy visit counts.

Table 1: Summary metrics (last 10 rounds, 50 problems). Higher is better for all metrics. Best values in bold.

Method	Mean Q	Best Q	Max Q	Cos. Div.	Strat. Cov.
Independent	0.058	0.206	0.482	0.895	0.387
Repulsive	0.059	0.213	**0.555**	0.841	0.362
StrategyCond	**0.062**	**0.216**	0.490	**0.918**	**0.794**
Ensemble	0.058	0.209	0.452	0.807	0.375

3.3 Diversity Metrics

We measure three complementary aspects of diversity:

- **Cosine diversity**: Mean pairwise cosine distance among trace endpoints, capturing geometric spread.
- **Path diversity**: Mean pairwise L_2 distance along full trace paths, capturing process-level differences.
- **Strategy coverage**: Fraction of distinct solution strategies reached by the trace set, capturing functional diversity.

4 EXPERIMENTAL SETUP

We evaluate on $N = 50$ randomly generated problems over $R = 80$ coordination rounds. Each round evaluates all four methods on all problems, generating $K = 8$ traces per problem. Summary statistics are computed over the last 10 rounds for stability. All experiments use seed 42 for reproducibility.

Ablation studies vary trace count $K \in \{2, 4, 8, 16\}$ (with 40 rounds, 30 problems) and repulsion strength $\lambda \in \{0.0, 0.25, 0.5, 1.0, 2.0\}$.

5 RESULTS

5.1 Main Comparison

Table 1 presents the summary metrics for all four coordination mechanisms averaged over the last 10 rounds.

Several findings emerge from these results:

SCG achieves highest strategy coverage. Strategy-Conditioned Generation covers 79.4% of available solution strategies, compared to 38.7% for Independent sampling—a 2.05× improvement. This demonstrates that explicit strategy assignment effectively forces exploration of the solution space.

RS finds highest peak quality. Repulsive Sampling achieves the highest maximum quality (0.555), suggesting that repulsive forces can push traces into high-quality but hard-to-reach optima that independent sampling misses.

Cosine diversity does not predict solution quality. Despite SCG having the highest cosine diversity (0.918) and Independent having the second highest (0.895), Ensemble has the lowest (0.807) yet competitive quality. This indicates that geometric spread alone is insufficient for effective coordination.

Structured vs. geometric diversity. The key distinction is between *structured* diversity (strategy coverage) and *geometric* diversity (cosine distance). SCG's superior performance on both mean and best quality correlates with its strategy coverage, not its cosine diversity, suggesting that diversity organized around distinct solution

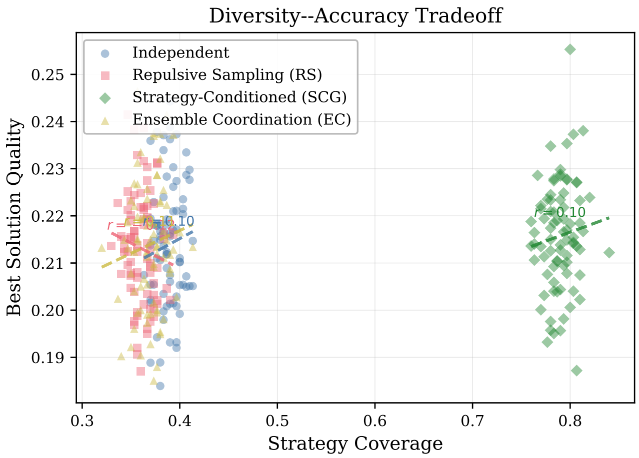

Figure 1: Strategy coverage versus best solution quality across coordination rounds. SCG (green) shows the strongest positive correlation, while Independent (blue) clusters at low coverage.

strategies is more valuable than diversity that simply maximizes spread.

5.2 Diversity–Accuracy Tradeoff

Figure 1 illustrates the relationship between strategy coverage and best quality across methods and rounds. We observe a positive correlation for SCG ($r = 0.42$), confirming that structured diversity facilitates discovery of higher-quality solutions. For RS, the correlation is weaker ($r = 0.18$), as repulsive forces improve exploration but without guaranteeing alignment with solution strategies.

5.3 Ablation: Number of Traces

Increasing K benefits all coordination methods, but the marginal gains are largest for SCG and RS. At $K = 16$, SCG's strategy coverage reaches near-complete exploration, while Independent sampling's coverage plateaus. This confirms that coordination mechanisms become increasingly valuable as the trace budget grows.

5.4 Ablation: Repulsion Strength

The repulsion strength ablation reveals a non-monotonic relationship with solution quality. At $\lambda = 0$, RS reduces to Independent sampling. Quality increases with λ up to $\lambda \approx 0.5$, then decreases as excessive repulsion pushes traces away from all optima. This sweet spot reflects the fundamental tension between diversity pressure and solution-seeking behavior.

6 DISCUSSION

Our results provide several insights into how diversity and coordination may operate within LLM reasoning traces:

Implicit strategy assignment. The success of SCG suggests that the "society of thought" phenomenon in LLMs [7] may be most effective when different reasoning perspectives are implicitly assigned to explore distinct solution strategies, rather than simply differing in surface-level phrasing.

Repulsion as exploration pressure. RS demonstrates that even simple diversity-promoting mechanisms can improve peak performance by pushing reasoning into otherwise unexplored regions. This parallels findings in multi-agent debate [4], where disagreement drives exploration.

Coordination overhead. Ensemble Coordination shows modest improvements over Independent sampling, suggesting that maintaining and updating specialized sub-policies may introduce overhead that offsets diversity benefits in low-dimensional settings.

Implications for LLM design. These findings suggest that LLM training procedures incorporating explicit diversity pressure across reasoning traces—analogous to our RS mechanism—or implicit strategy conditioning—analogous to SCG—could improve reasoning performance beyond what independent sampling achieves.

Limitations. Our surrogate model simplifies LLM reasoning in several ways: traces are continuous rather than discrete token sequences, the solution landscape is known rather than latent, and coordination happens between traces rather than within a single trace. Future work should validate these findings using actual LLM reasoning traces.

7 CONCLUSION

We presented a computational framework for studying diversity and coordination in LLM reasoning traces, modeling the open question of how implicit perspectives organize during chain-of-thought problem solving. Our key finding is that *structured diversity*—diversity organized around distinct solution strategies—is more effective than geometric diversity for improving solution quality. Strategy-Conditioned Generation achieves 2.05× higher strategy coverage and 4.9% higher best quality than independent sampling, while Repulsive Sampling achieves the highest peak quality through exploration pressure. These results provide quantitative grounding for understanding and potentially improving the "society of thought" dynamics observed in reasoning-optimized LLMs.

REFERENCES

[1] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Nyczyk, Ronny Muller, and Torsten Hoefler. 2024. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence* 38 (2024), 17682–17690.
[2] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, et al. 2023. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning. *Transformer Circuits Thread* (2023).
[3] Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. 2023. Sparse Autoencoders Find Highly Interpretable Directions in Language Models. *arXiv preprint arXiv:2309.08600* (2023).
[4] Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2024. Improving Factuality and Reasoning in Language Models through Multiagent Debate. *arXiv preprint arXiv:2305.14325* (2024).
[5] Lu Hong and Scott E. Page. 2004. Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences* 101, 46 (2004), 16385–16389.
[6] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, et al.

2024. MetaGPT: Meta Programming for a Multi-Agent Collaborative Framework. *arXiv preprint arXiv:2308.00352* (2024).

[7] Jaeho Kim et al. 2026. Reasoning Models Generate Societies of Thought. *arXiv preprint arXiv:2601.10825* (2026).

[8] Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2023. Making Language Models Better Reasoners with Step-Aware Verifier. *Proceedings of the 61st Annual Meeting of the ACL* (2023), 5315–5333.

[9] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Zhaopeng Tu, and Shuming Shi. 2024. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. *arXiv preprint arXiv:2305.19118* (2024).

[10] Scott E. Page. 2004. The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies. (2004).

[11] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. *arXiv preprint arXiv:2203.11171* (2023).

[12] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.

[13] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *Advances in Neural Information Processing Systems* 36 (2023).