Role-Based Orchestration of sLLM Agents for Korean Instruction-Following: A Comparison with SOTA

Extended Abstract

Gukhwan Bae
Daegu National University of Education
Daegu, South Korea
qo1598@dnue.ac.kr

Youngho Lee
Daegu National University of Education
Daegu, South Korea
yhlee@dnue.ac.kr

ABSTRACT

Motivated by the scarcity of on-shore GPUs and curated Korean corpora, we investigate whether Role-Based Orchestration (RBO), grounded in Role-Based Collaboration (RBC) theory, can serve as an effective alternative to single large models for Korean instruction-following tasks. We implement three agents—Generator, Critiquer, and Reviser—endowed with explicit responsibilities within an RBO framework that utilizes a rule-based Adaptive Controller and a Validator. Using 350 KoAlpaca items, we conduct evaluations comparing RBO against SOTA models like GPT-4o. While RBO achieves a competitive win rate in constraint-heavy tasks, its primary advantage lies in efficiency, requiring only $\approx$ 10% of the FLOPs per request compared to SOTA. This suggests a pragmatic path toward Sovereign AI in resource-constrained, non-English contexts.

KEYWORDS

Role-Based Orchestration, Multi-Agent Workflow, Small Language Models, Sovereign AI, LLM-as-a-Judge, Korean instruction-following

ACM Reference Format:
Gukhwan Bae and Youngho Lee. 2026. Role-Based Orchestration of sLLM Agents for Korean Instruction-Following: A Comparison with SOTA: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/GGVB5420

1 INTRODUCTION

LLM have achieved remarkable performance due to scaling laws[7], but approaches relying on single massive models face limitations such as diminishing returns relative to hardware scaling, depletion of high-quality data, and massive power consumption[5]. Furthermore, English-centric large models often struggle to reflect linguistic and cultural particularities of non-English languages like Korean[9]. This underscores the need for region-specific approaches from a Sovereign AI perspective[4].

To address these challenges, we propose the **Role-Based Orchestration (RBO)** framework. RBO reinterprets the traditional Role-Based Collaboration and E-CARGO framework[11–15] for the LLM environment. According to RBC research, a role is defined as a set of explicit Responsibilities and Rights, aimed at enhancing productivity by reducing conflict and ambiguity[1–3]. Unlike existing Mixture-of-Agents approaches that focus on parallel diversity[6, 10], RBO assigns explicit roles (Generator, Critiquer, Reviser) and improves output quality through a structured sequential workflow. In this study, we verify whether RBO, orchestrating small Korean LLMs[8], can serve as a pragmatic, cost-effective alternative to SOTA models in resource-constrained environments.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/GGVB5420

2 THE RBO FRAMEWORK ARCHITECTURE

The RBO framework operates as a multi-agent system following a sequential critique-refinement workflow. As illustrated in Figure 1, the architecture consists of four key components tailored to ensure instruction fidelity.

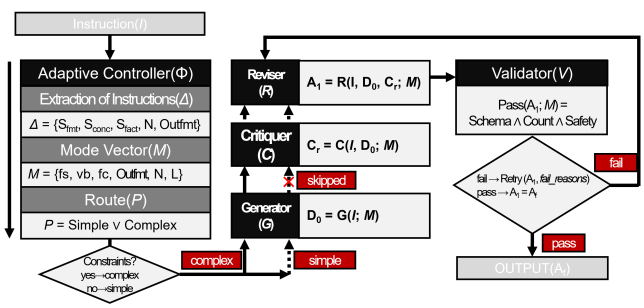

Figure 1: The overall architecture of the RBO framework. The Adaptive Controller routes the instruction, and the Validator ensures final constraint compliance.

2.1 Components

- **Adaptive Controller (Φ):** This module analyzes the user's instruction (I) to extract constraints such as format (S_{fmt}), count (S_{conc}), and factuality (S_{fact}). Based on these signals, it determines a *Mode Vector* (M) and selects an execution path (P). If the task requires strict adherence to constraints, the *Complex* path ($G \rightarrow C \rightarrow R$) is activated; otherwise, the *Simple* path is used.
- **Generator (G):** The Generator creates a structured draft (D_0) adhering to the instruction. To minimize hallucinations, its authority is strictly limited to the provided instruction and internal knowledge.
- **Critiquer (C):** This agent verifies the draft (D_0) against the Mode Vector (M). Instead of modifying the text, it produces a structured defect report (e.g., missing items, style violations), ensuring objective evaluation [15].

- **Reviser (R):** Taking the draft and the critique report, the Reviser generates an improved candidate output (A_1). Its goal is to maximize quality while strictly adhering to the original scope.
- **Validator (V):** As the final gatekeeper, the Validator performs a binary check (Pass/Fail) on the output based on explicit rules (Schema, Count, Safety). If validation fails, a retry loop is triggered to correct specific errors.

3 EXPERIMENTS AND RESULTS

3.1 Experimental Setup

We utilized the **KoAlpaca** dataset, a representative Korean instruction-following benchmark. A total of 350 items were sampled across 7 categories (Brainstorming, Classification, Closed QA, Generation, Information Extraction, Open QA, Summarization).

To strictly evaluate the efficacy of the role-based mechanism using sLLMs, we assigned specific Korean-centric models to each agent based on their capabilities: **Generator** (SKT A.X 7B), **Critiquer** (KT Mi:dm 2.3B), and **Reviser** (NAVER HyperCLOVA X 3B). We compared this RBO configuration against two SOTA models: **GPT-4o** and **DeepSeek-V3.2**. Performance was evaluated using a multi-model LLM-as-a-Judge panel (Claude-4.0-Sonnet, Gemini-2.5-Pro, Grok-4) and blinded human evaluation.

3.2 Performance Analysis

Overall, RBO achieved a win rate of 38.8% against GPT-4o and 43.1% against DeepSeek-V3.2. We verified the robustness of these results through high test-retest reproducibility (Pearson $r \approx 0.96$) across independent runs.

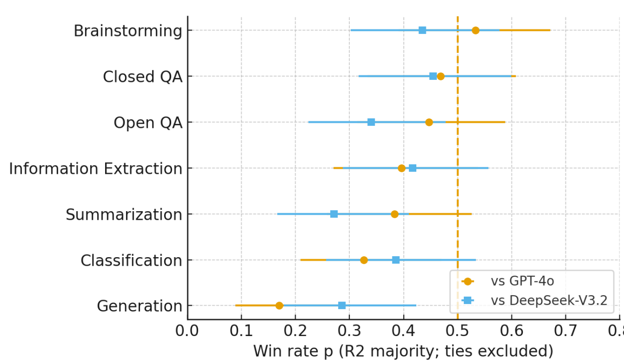

Figure 2: Win rates of RBO against SOTA models by task category.

Qualitative Analysis : As shown in Figure 2, RBO exhibited distinct performance patterns based on task types.

- **Strength (Brainstorming):** RBO outperformed GPT-4o in Brainstorming tasks ($p = 0.533$). Qualitative analysis of the judges' rationales reveals that the Validator's strict enforcement of 'Count' and 'Format' rules was the decisive factor. For example, in tasks requiring '5 items,' RBO consistently met the requirement, whereas SOTA models occasionally failed to adhere to the exact count.
- **Weakness (Generation):** Conversely, RBO underperformed in Generation tasks ($p = 0.170$). The Critiquer's conservative feedback tends to suppress creative expression, leading to outputs described by judges as 'dry and resume-like.' This highlights a trade-off between safety/strictness and creative diversity in the current rule-based setup.

3.3 Efficiency Analysis

The most significant advantage of RBO lies in its efficiency. As shown in Figure 3, RBO positions itself as an efficiency-centric model.

- **Computational Load:** RBO requires approximately **1.82 TFLOPs** per request, which is roughly **1/10th** of the computational load of DeepSeek-V3.2 ($\approx$ 17.45 TFLOPs).
- **Cost Effectiveness:** Operating with on-premise sLLMs, RBO's API cost is effectively $\approx$ **$0**, whereas GPT-4o incurs an estimated cost of $0.46 per comparable workload.
- **Latency Trade-off:** While RBO's E2E latency ($\approx$ 6.14s) is higher than GPT-4o ($\approx$ 2.56s) due to sequential processing, this is an acceptable trade-off in scenarios where data privacy and operational costs are the primary constraints.

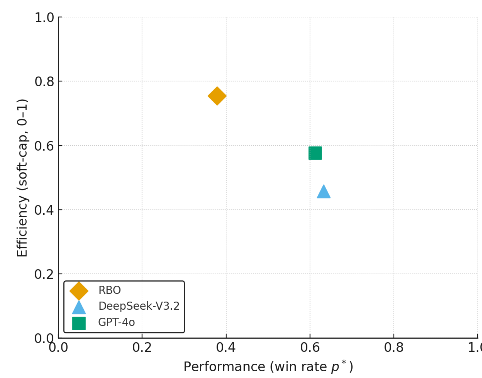

Figure 3: Performance-Efficiency Trade-off. RBO demonstrates superior efficiency with competitive performance compared to SOTA models.

4 CONCLUSION

This study demonstrates that the Role-Based Orchestration of small, Korean-centric LLMs is a viable alternative to SOTA models for Korean instruction-following tasks. By leveraging explicit role definitions and a validation mechanism, RBO ensures high fidelity to constraints. Most importantly, it achieves this at a fraction of the compute and cost, aligning with Sovereign AI objectives where on-premise deployment and accessibility are paramount. Future work will explore learning-based routing to further mitigate the trade-off between creativity and strictness.

5 ACKNOWLEDGMENTS

This work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2025S1A5A8006876)

REFERENCES

[1] 2022. Tackling Challenges of Robustness Measures for Autonomous Agent Collaboration in Open Multi-Agent System. In *Hawaii International Conference on System Sciences*. IEEE Computer Society.

[2] 2022. Unrolled Graph Learning for Multi-Agent Collaboration. arXiv:2210.17101 [cs.MA] https://arxiv.org/abs/2210.17101

[3] Abdelghani Boudjidj and Mohammed El Habib Souidi. 2022. A Formal Composition of Multi-Agent Organization based on Category Theory. In *Computer Science & Information Technology*. Academy and Industry Research Collaboration Center. https://doi.org/10.5121/csit.2022.121909

[4] Alexandre Costa-Barbosa, Bianca Herlo, Gesche Joost, et al. 2024. Digital Sovereignty in times of AI: between perils of hegemonic agendas and possibilities of alternative approaches. *Liinc em Revista* 20, 2 (2024).

[5] Jared Fernandez, Luca Wehrstedt, Leonid Shamis, Mostafa Elhoushi, Kalyan Saladi, Yonatan Bisk, Emma Strubell, and Jacob Kahn. 2025. Hardware Scaling Trends and Diminishing Returns in Large-Scale Distributed Training. arXiv:2411.13055 [cs.LG] https://arxiv.org/abs/2411.13055

[6] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of Experts. arXiv:2401.04088 [cs.LG] https://arxiv.org/abs/2401.04088

[7] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling Laws for Neural Language Models. arXiv:2001.08361 [cs.LG] https://arxiv.org/abs/2001.08361

[8] Hyun-Jae Lee, Jaewoong Yoon, Bonggyu Hwang, Seongho Joe, Seungjai Min, and Youngjune Gwon. 2021. KoreALBERT: Pretraining a Lite BERT Model for Korean Language Understanding. arXiv:2101.11363 [cs.CL] https://arxiv.org/abs/2101.11363

[9] Aditya Siddhant, Junjie Hu, Melvin Johnson, Orhan Firat, and Sebastian Ruder. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization. In *Proceedings of the International Conference on Machine Learning*, Vol. 2020. 4411–4421.

[10] Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2024. Mixture-of-Agents Enhances Large Language Model Capabilities. arXiv:2406.04692 [cs.CL] https://arxiv.org/abs/2406.04692

[11] Shiyu Wu, Meiqiao Pan, Yanyan Fan, Shenglin Li, and Libo Zhang. 2023. Group role assignment with a training plan considering the duration in adaptive collaboration. In *2023 26th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. IEEE, 739–744.

[12] Hui Zhao, Wanlong Li, Xiuping Hou, and Shanhong Zheng. 2009. A Domain-Role-Based Control Model for CSCW System. In *2009 International Conference on Management and Service Science*. IEEE, 1–4.

[13] Haibin Zhu. 2006. A role-based approach to robot agent team design. In *2006 IEEE International Conference on Systems, Man and Cybernetics*, Vol. 6. IEEE, 4861–4866.

[14] Haibin Zhu. 2021. *E-CARGO and role-based collaboration: Modeling and solving problems in the complex world*. John Wiley & Sons.

[15] Haibin Zhu and MengChu Zhou. 2006. Role-based collaboration and its kernel mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 36, 4 (2006), 578–589.