

# Redefining Evaluation Standards: A Unified Framework for Evaluating the Korean Capabilities of Language Models

Hanwool Lee<sup>1,2,\*</sup>, Dasol Choi<sup>1,\*</sup>, Sooyong Kim<sup>3</sup>, Ilgun Jeong<sup>4</sup>, Sangwon Baek<sup>5</sup>,  
Guijin Son<sup>2</sup>, Inseon Hwang<sup>6</sup>, Naeun Lee<sup>7</sup>, Seunghyeok Hong<sup>8,†</sup>

<sup>1</sup>AIM Intelligence <sup>2</sup>Seoul National University <sup>3</sup>A.I.MATICS <sup>4</sup>TigerCompany <sup>5</sup>Catius

<sup>6</sup>National Assembly of Korea <sup>7</sup>Coupang <sup>8</sup>Hankuk University of Foreign Studies

gksdnf424@gmail.com

## Abstract

Recent advances in Korean large language models (LLMs) have produced numerous benchmarks and evaluation methods, yet inconsistent protocols lead to performance gaps of up to 10 percentage points across institutions. These discrepancies arise from fragmented practices in prompt templates, inference settings, and evaluation criteria, making reliable comparison difficult. We introduce **HRET (Haerae Evaluation Toolkit)**, an open-source, registry-based framework that unifies Korean LLM evaluation. HRET standardizes data ingestion, inference, and reporting pipelines to ensure reproducibility, while remaining flexible through modular integration of major Korean benchmarks, multiple inference backends, and diverse evaluation methods (string matching, log-likelihood, mathematical verification, and LLM-as-judge). Beyond standard accuracy metrics, HRET provides Korean-specific diagnostic analyses—morphology-aware Type–Token Ratio for lexical diversity and systematic keyword-omission detection for conceptual coverage—offering linguistic insights beyond conventional evaluation. Across diverse hardware and backend configurations, HRET yields score variations below 0.003 (vs. 0.1+ pp gaps reported previously), demonstrating its effectiveness in eliminating evaluation discrepancies and establishing a reproducible standard for Korean LLM benchmarking. HRET is publicly available at <https://github.com/HAE-RAE/haerae-evaluation-toolkit>.

**Keywords:** Korean NLP, Evaluation Framework, Reproducibility, Benchmarking, LLM-as-Judge

## 1. Introduction

Recent advances in large language models (LLMs) demand rigorous and reproducible evaluation, yet real-world benchmarking often diverges from official protocols (Liang et al., 2022; Gao et al., 2024). In Korean NLP, identical models have shown discrepancies of up to 10 percentage points across institutions (Research et al., 2024; Bak et al., 2025), largely due to variations in prompt templates, inference settings, and evaluation criteria. Such inconsistency undermines fair comparison and slows progress in the field.

The reproducibility crisis in Korean LLM evaluation arises from two main factors. First, practitioners face a fragmented landscape: even widely used benchmarks such as KMMLU (Son et al., 2024b) in *lm-eval-harness* are rarely applied without modification. Researchers frequently augment them with few-shot examples (Brown et al., 2020), chain-of-thought prompts, or self-consistency voting (Wang et al., 2023) to better capture reasoning. Second, Korean’s linguistic characteristics—rich morphology, flexible word order, and frequent case-marker omission—introduce tokenization and normalization differences that materially affect scores. Without standardized handling of these phenomena, reproducibility remains elusive.

While general-purpose frameworks such as *lm-*

*evaluation-harness* (Gao et al., 2024) and *DeepEval* (AI, 2025) have advanced English evaluation, they lack native support for Korean linguistic phenomena, requiring significant manual adaptation that often reintroduces the inconsistencies they aim to eliminate.

We introduce **HRET (Haerae Evaluation Toolkit)**, an open-source, registry-based framework that unifies Korean LLM assessment. HRET standardizes evaluation workflows—data ingestion, inference, and result reporting—while remaining flexible through modular integration of major Korean benchmarks, multiple inference backends, and diverse evaluation methods (string matching, log-likelihood, mathematical verification, and LLM-as-judge). It further enforces language consistency to penalize non-Korean outputs and incorporates Korean-specific diagnostics: morphology-aware Type–Token Ratio for lexical diversity and systematic keyword-omission detection for conceptual coverage. Beyond its pipeline-level design, HRET provides a high-level decorator-style API inspired by *DeepEval*, enabling one-line evaluation, multi-model benchmarking, and seamless MLOps tracking. We validate HRET across diverse hardware and backends, showing performance variations below 0.003 standard deviation—an order of magnitude smaller than prior reports—demonstrating a practical path toward reproducible, adaptable Korean LLM evaluation.

\*Equal contribution.

†Corresponding author.

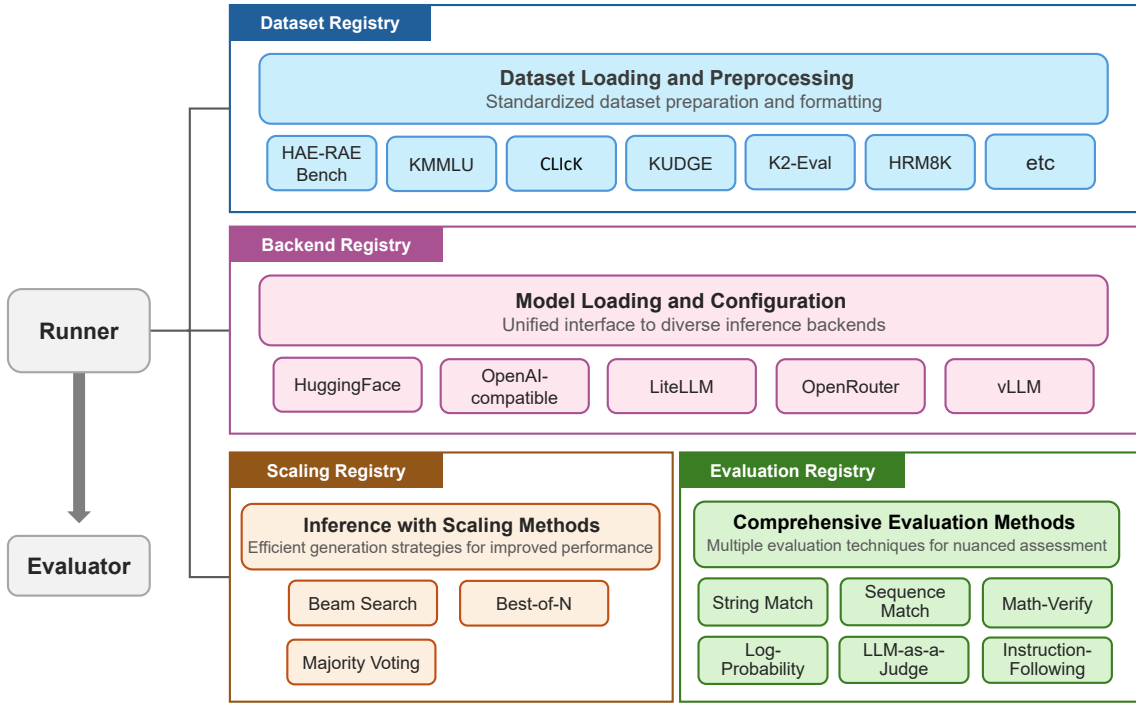


Figure 1: Architecture overview of HRET, showing the registry-based modular design that integrates datasets, models, inference strategies, and evaluation methods through a unified pipeline.

Report	Model	Datasets				
		GSM8K	MMLU	KMMLU	HumanEval	MBPP
EXAONE 3.5 Report (Research et al., 2024)	Qwen2.5-32B	92.00	81.40	62.10	89.00	88.90
	Gemma2-27B	84.20	74.80	53.80	79.30	80.70
Kanana Report (Bak et al., 2025)	Qwen2.5-32B	95.30	84.40	59.37	82.32	71.96
	Gemma2-27B	91.05	78.30	49.98	70.12	70.90

Table 1: Performance discrepancies for identical models between EXAONE 3.5 and Kanana technical reports, highlighting reproducibility challenges in Korean NLP evaluations.

## 2. Related Work and Background

### 2.1. Global LLM Benchmarks and Evaluation Frameworks

The English LLM ecosystem has established mature evaluation frameworks such as Evalverse (Kim et al., 2024b) and DeepEval (AI, 2025). Among them, the *lm-evaluation-harness* (Gao et al., 2024) has been particularly influential, powering the Open LLM Leaderboard and enabling transparent comparison across models. Its success demonstrates that accessible evaluation tools can democratize model development—allowing small research groups and startups to participate on equal footing. Similarly, the Chinese framework FlagEval (Team, 2025) has advanced its language ecosystem through an accessibility-first approach that integrates datasets, metrics, and automation into a unified hub.

### 2.2. Reproducibility Challenges in Korean LLM Evaluation

Evaluations of Korean language models show notable inconsistencies across institutions. Identical models often report performance gaps of up to ten percentage points on benchmarks such as GSM8K, KMMLU, HumanEval, and MBPP when comparing the EXAONE 3.5 report (Research et al., 2024) and the Kanana report (Bak et al., 2025) (Table 1). These discrepancies arise from variations in prompt templates, inference settings, and evaluation criteria, revealing the absence of standardized Korean-specific practices.

Beyond procedural factors, linguistic characteristics of Korean intensify these challenges. Its rich morphology, flexible word order, and frequent case-marker omission introduce tokenization and normalization differences that can shift results materially (Jeon et al., 2023; Kim et al., 2024c). Without consistent language-aware handling, results risk

reflecting preprocessing artifacts rather than true model competence.

### 2.3. Language-Aware Benchmarking for Korean

While reproducibility issues are procedural in part, they are amplified by the linguistic complexity of Korean, which English-oriented evaluation pipelines often fail to capture (Arnett and Bergen, 2024). As an agglutinative and morphologically rich language, Korean forms words through variable suffixation and inflection. Naive subword tokenization (e.g., BPE or WordPiece) frequently misaligns with true morpheme boundaries, reducing evaluation reliability. Moreover, flexible word order and frequent case-marker omission make reference matching ambiguous without morphological or syntactic normalization. Differences in tokenization or spacing can shift scores by several points, conflating model ability with preprocessing artifacts. Therefore, a language-aware evaluation framework—featuring morpheme alignment, spacing normalization, and honorific-aware matching—is essential to isolate linguistic competence from implementation variance.

### 2.4. Comparison with Existing Systems

Table 2 situates HRET within the broader landscape of evaluation frameworks. HRET differentiates itself through three design principles: (1) **Korean-specific capabilities** such as language-consistency enforcement and morphology-based diagnostics; (2) **registry-based extensibility**, enabling rapid integration of new datasets and methods without code restructuring; and (3) **a unified evaluation pipeline** that standardizes fragmented Korean LLM assessment practices. These choices collectively form the foundation of HRET, described in the following section, which operationalizes these ideas into a reproducible, modular evaluation framework.

## 3. HRET Framework

**HRET (Haerae Evaluation Toolkit)** addresses the reproducibility and comparability challenges inherent in evaluating Korean large language models (LLMs). It provides a standardized yet flexible environment that ensures consistency, transparency, and accessibility while accommodating diverse research and deployment needs. Figure 1 illustrates its registry-based architecture, which integrates datasets, models, inference strategies, and evaluation methods into a unified pipeline.

HRET targets three primary user groups: (1) researchers seeking reproducible Korean benchmarks; (2) developers requiring reliable evaluation

tools for model selection; and (3) practitioners deploying Korean LLMs in production environments.

### 3.1. Design Principles

**Modular and Scalable Architecture.** All components—dataset loaders, evaluators, scaling strategies, and backends—are registered modularly, allowing new elements to be added with minimal code modification. Parallelized execution using asynchronous I/O and multi-threading enables efficient large-scale benchmarking.

**Standardized and Reproducible Evaluation.** HRET unifies data ingestion, inference, and reporting through consistent interfaces and YAML-based configurations, ensuring identical results across hardware setups and institutions.

**Language-Aware Design.** The toolkit enforces Korean-only outputs via a language consistency penalty and integrates morphology- and semantics-aware diagnostics to evaluate linguistic robustness beyond surface accuracy.

### 3.2. System Architecture

HRET follows a standardized five-stage workflow—dataset loading, model setup, inference, evaluation, and reporting—covering major Korean benchmarks such as HAE-RAE, KMMLU, KUDGE, HRM8K, and CLiCK (see Section 3.3).

1. **Dataset Registry:** Provides consistent dataset formats and a generic loader for custom files.
2. **Backend Registry:** Connects to HuggingFace, vLLM, or OpenAI-compatible APIs through a unified interface that supports both generation and judging modes.
3. **Inference Scaling:** Includes beam search, best-of- $n$  sampling, and majority voting to stabilize outputs and capture diverse reasoning paths.
4. **Evaluation:** Supports four complementary paradigms—(i) **String Matching** for discrete QA and short-form tasks; (ii) **Log-Likelihood Scoring** for multiple-choice and completion-style benchmarks; (iii) **Mathematical Verification** for reasoning and equation-consistency checks; and (iv) **LLM-as-a-Judge** for open-ended or generative responses requiring qualitative assessment. Language-consistency enforcement penalizes non-Korean outputs, and automated answer extraction ensures reliable comparison across reasoning styles.
5. **Reporting:** Aggregates runtime outputs from all registries through a facade-style Evaluator API for unified logging, visualization, and downstream analysis.

Feature	lm-eval-harness (Biderman et al., 2024)	OpenAI Evals (Liu et al., 2024)	Evalverse (Kim et al., 2024b)	DeepEval (AI, 2025)	HRET (Ours)
<b>Primary Focus</b>	Academic benchmarks	OpenAI model evaluation	Unified evaluation hub	Unit/regression testing	Korean LLM evaluation
<b>Developer</b>	EleutherAI	OpenAI	Upstage	Confident AI	Community Contributors
<b>Architecture</b>	Task/Metric modules	Evaluation scripts	Submodule integration	Test case-based	Registry-based plugins
<b>Built-in Benchmarks</b>	60+ general benchmarks	Several public evals	Integrated external benchmarks	Major benchmarks	Korean-focused benchmarks
<b>Evaluation Metrics</b>	Traditional metrics	Task-specific scripts	Integrated metrics	5+ diverse metrics	5+ diverse metrics
<b>User Interface</b>	CLI only	CLI + Dashboard	Slack bot + Web	CLI + Web dashboard	CLI + Python API
<b>Automation</b>	Script-based	API-based	No-code automation	CI/CD integration	API-based, CI/CD integration
<b>Unique Strength</b>	High reproducibility	OpenAI optimization	Multi-tool integration	Testing-focused	Korean-specific + Registry extensibility

Table 2: Comparison of major evaluation frameworks. HRET uniquely provides registry-based extensibility and Korean-specific benchmarks.

Category	Benchmark	Description
<b>General</b>	HAE-RAE	Evaluates general Korean language understanding, focusing on vocabulary, reading comprehension, and broad knowledge.
	KMMLU	Covers 45 knowledge categories (STEM, humanities, applied sciences) in Korean.
	CLlck	Tests cultural and linguistic proficiency using 1,995 QA pairs from Korean exams and textbooks.
<b>Dialogue</b>	KUDGE	Assesses coherence and contextual appropriateness in Korean dialogue generation.
<b>Specialized</b>	K2-Eval	Tests professional and technical Korean terminology.
	HRM8K	Evaluates mathematical reasoning across 8,011 Korean problems.
	HRMCR	Multi-step commonsense reasoning using Korean cultural templates.
	HRC	Korean reasoning challenge with 119 science, mathematics, and quiz tasks.
	KorMedQA	7,469 Korean medical licensing MCQs (2012–2024).
	KBL	Legal reasoning and bar-exam questions in closed-book and RAG settings.
	KoBalt	Evaluates linguistic competence across syntax, semantics, pragmatics, phonology, and morphology via 700 expert-curated items.
	K-Halu	Tests hallucination detection across seven knowledge domains from Korean media sources.
	Ko-IFEval	Korean adaptation of Google’s IFEval dataset for instruction-following assessment.
Benchhub	Comprehensive evaluation suite integrating multiple Korean instruction-following and reasoning benchmarks under unified schemas.	

Table 3: Benchmarks currently supported in HRET. Datasets span general, dialogue, and specialized Korean language evaluations.

### 3.3. Supported Benchmarks

HRET provides unified evaluation support for major Korean language benchmarks spanning diverse domains and task types. Its registry-based architecture enables seamless integration of these benchmarks with consistent preprocessing, evaluation protocols, and result formatting.

The collection covers three major categories: **(1) General** benchmarks such as HAE-RAE (Son et al., 2023), KMMLU (Son et al., 2024b), and CLlck (Kim et al., 2024a), which assess broad Korean language understanding and cultural knowledge; **(2) Dialogue** benchmarks like KUDGE (Son et al., 2024a), which evaluate conversational coherence and contextual appropriateness; and **(3) Specialized** benchmarks that target domain-specific reasoning and professional knowledge, including K2-Eval (Son et al., 2024a), HRM8K (Ko et al., 2025), HRMCR (Son et al., 2025), KorMedQA (Kweon

et al., 2024), KBL (Kim et al., 2024d), KoBalt (Shin et al., 2025), K-Halu (Seo and Lim, 2025), Ko-IFEval (Lee and Team, 2024), and Benchhub (Kim et al., 2025). This comprehensive coverage enables holistic evaluation across linguistic, reasoning, and professional knowledge domains.

Table 3 provides detailed descriptions of the supported benchmarks organized by evaluation focus.

### 3.4. Advanced and Extensible Pipelines

Beyond standard evaluation, HRET’s PipelineRunner enables advanced workflows including multi-model comparisons, ensemble scoring, and cascading evaluations where one model’s outputs feed another. A high-level decorator interface inspired by DeepEval simplifies experiment definition while maintaining compatibility with the underlying pipeline:

```
import llm_eval.hret as hret
```

```
@hret.evaluate(dataset="kmmlu",
              model="huggingface",
              evaluation_method
                ="string_match")
def my_model(x: str) -> str:
    return model.generate(x)
result = my_model()
print(result.metrics["accuracy"])
```

Both interfaces are interoperable and integrate seamlessly with MLOps tools (e.g., MLflow, W&B) for experiment tracking and result persistence.

### 3.5. Korean-Specific Diagnostic Analyses

HRET extends beyond conventional accuracy metrics by incorporating morphology-aware analyzers and semantic diagnostics to capture language-specific phenomena. It provides four diagnostic views: (1) **Performance Overview** with detailed accuracy and prediction distributions; (2) **Subset Analysis** across domains (e.g., reasoning, culture, comprehension); (3) **Lexical Diversity** via Type-Token Ratio (TTR) between correct and incorrect responses; and (4) **Error Pattern Detection** through systematic keyword omission tracking. These analyses expose morphological and semantic weaknesses invisible to surface metrics, guiding targeted model refinement.

HRET also defines three extended task families—(i) instruction following (schema adherence and refusal compliance), (ii) long-context reasoning (cross-chunk consistency), and (iii) hallucination stress tests (factuality under distractors). Each family includes metric adapters and Korean honorific-aware judge prompts to ensure culturally and linguistically appropriate assessment.

### 3.6. Usage and Reproducibility

HRET supports both command-line and Python interfaces for reproducible benchmarking. A minimal evaluation example is shown below, with extended configurations and templates available in the project repository.

```
python -m llm_eval.evaluator \
  --model huggingface \
  --dataset haerae_bench \
  --evaluation_method
  string_match
```

## 4. System Evaluation

### 4.1. Reproducibility Validation

To evaluate HRET’s effectiveness in addressing reproducibility challenges, we conducted cross-institutional experiments across multiple hardware configurations and backend implementations. Three representative models—Qwen2.5-3B (Qwen et al., 2025), Qwen3-8B (Yang et al., 2025), and Llama-3.2-3B-Instruct (Meta AI, 2024)—were tested on Korean benchmarks using log-likelihood evaluation across four heterogeneous environments differing in CPU architecture, GPU type, and memory capacity.

Model	Max Variation	Hardware Configs
Qwen2.5-3B	0.0016 (0.56%)	4 setups
Qwen3-8B	0.0028 (1.06%)	4 setups
Llama-3.2-3B	0.0015 (0.62%)	4 setups

Table 4: Hardware robustness summary across environments.

Table 5 provides the full per-configuration results. Across server-grade (A100, H100) and workstation-grade (RTX A6000, RTX 3090) environments, scores remain nearly identical, confirming HRET’s hardware-agnostic consistency.

These results confirm that HRET effectively eliminates the 1–10 percentage point discrepancies previously observed between institutions (cf. Table 1). Within each backend, results remain statistically identical ( $\sigma < 0.003$ ), confirming that HRET ensures reproducible evaluation even under varied computational environments.

### 4.2. Config-Driven Reproducibility via YAML

To facilitate exact replication across institutions, HRET employs standardized, shareable YAML configurations. For all integrated benchmarks, published scores can be reproduced with less than 1% deviation by matching only configuration keys—such as prompt templates, decoding parameters, and judge settings—while keeping dataset splits fixed:

```
# hret_config.yaml
default_dataset: "kmmlu"
default_model: "huggingface"
default_split: "test"
default_evaluation_method: "
  string_match"
batch_size: 32
max_workers: 4
```

```
import llm_eval.hret as hret
```

Model	Hardware Setup	HAE-RAE	CLlck	Batch Size
Qwen2.5-3B	AMD EPYC 74F3 + A100×2	0.2847	0.3388	64
	AMD EPYC 9354 + RTX A6000	0.2847	0.3388	2
	Intel i9-13900K + RTX A6000	0.2847	0.3388	2
	Intel Xeon + A100-40GB	0.2863	0.3378	64
Qwen3-8B	AMD EPYC 74F3 + A100×2	0.2645	0.3083	64
	AMD EPYC 9354 + RTX A6000	0.2645	0.3083	1–2
	Intel i9-13900K + RTX A6000	0.2645	0.3083	2
	Intel Xeon + A100-40GB	0.2673	0.3083	64
Llama-3.2-3B	AMD EPYC 74F3 + A100×2	0.2436	0.2922	64
	AMD EPYC 7452 + RTX 3090	0.2437	0.2922	2
	Intel i9-13900K + RTX A6000	0.2445	0.2922	2
	Intel Xeon + A100-40GB	0.2447	0.2907	64

Table 5: Reproducibility validation across hardware configurations using HuggingFace backend.

Hardware Configuration	HAE-RAE	CLlck
AMD EPYC 74F3 + H100	0.1806	0.1343
Intel Xeon + A100-80GB	0.1788	0.1298
Intel Xeon + H200	0.1802	0.1308
<b>Standard Deviation</b>	<b>0.0009</b>	<b>0.0023</b>

Table 6: Backend robustness using vLLM for Qwen2.5-3B.

```
hret.load_config("hret_config.
yaml")
res = hret.quick_eval(
my_model_function, dataset="
knmlu")
```

This configuration discipline eliminates hidden variability, enabling reproducible “paper-accurate” replication across teams and hardware.

### 4.3. Multi-Method Evaluation Analysis

Korean tasks vary widely in linguistic and reasoning complexity, requiring different evaluation strategies. HRET’s multi-method framework enables appropriate metric selection per task, balancing accuracy and interpretability.

Table 7 compares evaluation methods for EXAONE-3.5-2.4B (Research et al., 2024) across two benchmarks. Results reveal method sensitivity in reasoning tasks and stability in linguistic ones.

Benchmark	String-Match	Math-Verify	LLM-as-judge
HRM8K	2.21	14.11	17.60
CLlck	27.31	28.67	27.76

Table 7: Comparison of evaluation methods for EXAONE-3.5-2.4B using GPT-4o (2024-11-20) as the LLM-as-judge.

Mathematical reasoning (HRM8K) benefits sub-

stantially from semantic evaluation—scores increase from 2.2% (string-match) to 17.6% (LLM-as-judge)—whereas linguistic and cultural QA (CLlck) remain robust across metrics (27–29%). These findings validate HRET’s ability to align evaluation method with task type.

### 4.4. Test-Time Scaling Techniques

To further enhance evaluation robustness, HRET integrates test-time scaling strategies that average or aggregate predictions from multiple inference runs. The framework supports three methods: **Best-of-N Sampling**, **Self-Consistency** via majority voting, and **Beam Search** for structured decoding. These can be combined with any evaluation metric, including LLM-as-a-judge assessments.

Model	Benchmark	n=1	n=3	n=5
EXAONE-3.5-2.4B	HRC	3.37	5.04	9.24
Qwen2.5-3B-Instruct		11.76	12.60	13.44

Table 8: Test-time scaling with LLM-as-a-judge (GPT-4o-2024-11-20) on HRC.

Increased sampling significantly improves performance—EXAONE-3.5-2.4B gains  $2.7\times$  at  $n = 5$  compared to single-shot inference—confirming the benefit of stochastic aggregation for reasoning-intensive Korean tasks.

### 4.5. Performance Comparison with Existing Tools

Table 9 compares HRET’s computational performance with lm-evaluation-harness using Qwen2.5-3B on identical hardware configurations.

HRET shows mixed performance: faster on HAE-RAE Bench but slower on KMMLU. Higher RAM usage stems from loading datasets as editable objects to enable real-time filtering and experimental

Metric	HRET	lm-eval-harness
<b>HAE-RAE Bench</b>		
Execution Time (sec)	114.93	142.44
Max RAM (MB)	80.96 (<1%)	54.26 (<1%)
Max CPU (%)	525.84	100.0
<b>KMMLU</b>		
Execution Time (sec)	303.93	221.04
Max RAM (MB)	104.81 (<1%)	78.76 (<1%)
Max CPU (%)	458.90	199.0

Table 9: Performance comparison on AMD EPYC 74F3 64-Core + NVIDIA H100 with 193GB RAM. Values averaged across 3 runs.

modifications through our system. Higher CPU utilization (up to 525%) reflects multi-threading optimization. Both frameworks use <1% of total system memory, indicating efficient resource utilization despite HRET’s flexibility-oriented design.

Across all experiments, HRET demonstrates (i) near-zero variance across hardware and backends, (ii) configuration-based reproducibility within 1%, and (iii) flexible adaptation to diverse evaluation paradigms. Together, these results validate HRET as a robust, extensible foundation for reproducible and language-aware Korean LLM benchmarking.

#### 4.6. Extensibility

HRET’s registry-based design ensures seamless extensibility. New components—**Models**, **Datasets**, **Evaluators**, and **Scaling Methods**—can be added through a single decorator (e.g., `@register_model`), automatically integrating with the existing runtime without additional configuration.

This plug-and-play approach encourages community contributions and supports continuous evolution of the framework alongside emerging Korean benchmarks and inference paradigms, while maintaining full backward compatibility.

### 5. Real-World Adoption

Beyond academic evaluation, HRET has been adopted in production environments across the Korean NLP ecosystem. The Horangi Korean LLM Leaderboard<sup>1</sup>, a community-driven leaderboard hosted on Weights & Biases, uses HRET as its evaluation backbone to ensure standardized and reproducible model comparison. Additionally, the Haerae Reasoning Challenge (HRC) officially employs HRET to score all participant submissions

<sup>1</sup><https://wandb.ai/wandb-korea/korean-llm-leaderboard/reports/Horangi-LLM---Vmlldzo3MzIyNDE2>

under standardized protocols, ensuring fairness and transparency in competitive evaluations. Authors’ affiliated organizations, including OneLineAI, have also integrated HRET into their model validation and post-training evaluation workflows, demonstrating the framework’s practical utility in deployment scenarios. This real-world adoption validates HRET’s reliability beyond controlled research settings, confirming its value as evaluation infrastructure for both the research community and industry practitioners.

## 6. Conclusion

We introduce HRET, a unified evaluation framework that addresses critical reproducibility challenges in Korean LLM evaluation. Through its registry-based architecture and standardized protocols, HRET significantly improves evaluation consistency while supporting diverse assessment strategies. HRET integrates major Korean benchmarks with multiple inference backends and evaluation methods, enabling researchers to select task-appropriate strategies while maintaining reproducibility. The framework’s extensible design and community-driven development ensure it evolves with the Korean LLM landscape, providing a sustainable foundation for reproducible assessment.

## Limitations

While HRET provides a standardized and reproducible evaluation framework for Korean LLMs, several limitations remain. First, although it integrates most major Korean benchmarks, coverage in specialized domains (Hwang et al., 2025; Choi et al., 2025) is still limited, and the current framework does not yet support vision-language model (VLM) benchmarks (Choi et al., 2026), which are increasingly important for multimodal Korean evaluation. Second, the language consistency enforcement may struggle with mixed-language or code-switched outputs common in technical and multilingual contexts. Third, test-time scaling techniques require considerable computational resources, which may limit accessibility for smaller institutions or resource-constrained environments. In addition, current metrics—particularly string matching—do not fully capture Korean’s morphological richness and honorific nuances, and the framework’s accuracy ultimately depends on the quality of integrated benchmarks. We are actively addressing these challenges through ongoing development and invite community contributions to further enhance HRET’s scope and robustness.

## Ethics Statement

HRET aims to improve reproducibility and fairness in Korean LLM evaluation. However, we acknowledge several ethical considerations: (1) standardized evaluation frameworks may inadvertently favor certain model architectures or training approaches; (2) benchmark selection reflects our research perspectives and may not capture all aspects of Korean language competency; (3) language consistency enforcement, while promoting Korean language assessment, may disadvantage multilingual models in certain contexts. We encourage the community to contribute diverse benchmarks and evaluation methods to ensure broad representation of Korean language capabilities. All integrated datasets respect original licensing and guidelines.

## 7. References

- Confident AI. 2025. [Deepeval: The llm evaluation framework](#). Accessed: 2025-01-26.
- Catherine Arnett and Benjamin K. Bergen. 2024. [Why do language models perform worse for morphologically complex languages?](#)
- Yunju Bak, Hojin Lee, Minho Ryu, Jiyeon Ham, Seungjae Jung, Daniel Wontae Nam, Taegyong Eo, Donghun Lee, Dooha Jung, Boseop Kim, et al. 2025. [Kanana: Compute-efficient bilingual language models](#). *arXiv preprint arXiv:2502.18934*.
- Stella Biderman, Hailey Schoelkopf, Lintang Sutawika, Leo Gao, Jonathan Tow, Baber Abbasi, Alham Fikri Aji, Pawan Sasanka Ammanamanchi, Sidney Black, Jordan Clive, et al. 2024. [Lessons from the trenches on reproducible evaluation of language models](#). *arXiv preprint arXiv:2405.14782*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.
- Dasol Choi, Jungwhan Kim, and Guijin Son. 2025. [Ko-piqa: A korean physical commonsense reasoning dataset with cultural context](#). *arXiv preprint arXiv:2509.11303*.
- Dasol Choi, Guijin Son, Hanwool Lee, Minhyuk Kim, Hyunwoo Ko, Teabin Lim, Ahn Eungyeol, Jungwhan Kim, Seunghyeok Hong, and Youngsook Song. 2026. [What users leave unsaid: Under-specified queries limit vision-language models](#). *arXiv preprint arXiv:2601.06165*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2024. [A framework for few-shot language model evaluation](#).
- Yewon Hwang, Sungbum Jung, Hanwool Lee, and Sara Yu. 2025. [Twice: What advantages can low-resource domain-specific embedding model bring? – a case study on korea financial texts](#).
- Taehee Jeon, Bongseok Yang, Changhwan Kim, and Yoonseob Lim. 2023. [Improving korean nlp tasks with linguistically informed subword tokenization and sub-character decomposition](#).
- Eunsu Kim, Juyoung Suk, Philhoon Oh, Haneul Yoo, James Thorne, and Alice Oh. 2024a. [Click: A benchmark dataset of cultural and linguistic intelligence in korean](#). *arXiv preprint arXiv:2403.06412*.
- Eunsu Kim, Haneul Yoo, Guijin Son, Hitesh Patel, Amit Agarwal, and Alice Oh. 2025. [Benchhub: A unified benchmark suite for holistic and customizable llm evaluation](#).
- Jihoo Kim, Wonho Song, Dahyun Kim, Yunsu Kim, Yungi Kim, and Chanjun Park. 2024b. [Evalverse: Unified and accessible library for large language model evaluation](#). *arXiv preprint arXiv:2404.00943*.
- Jong Myoung Kim, Young-Jun Lee, Yong jin Han, Sangkeun Jung, and Ho-Jin Choi. 2024c. [Does incomplete syntax influence korean language model? focusing on word order and case markers](#).
- Yeeun Kim, Young Rok Choi, Eunkyung Choi, Jinhwan Choi, Hai Jin Park, and Wonseok Hwang. 2024d. [Developing a pragmatic benchmark for assessing korean legal language understanding in large language models](#).
- Hyunwoo Ko, Guijin Son, and Dasol Choi. 2025. [Understand, solve and translate: Bridging the multilingual mathematical reasoning gap](#). *arXiv preprint arXiv:2501.02448*.
- Sunjun Kweon, Byungjin Choi, Gyok Chu, Junyeong Song, Daeun Hyeon, Sujin Gan, Jueon Kim, Minkyu Kim, Rae Woong Park, and Edward

- Choi. 2024. [Kormedmcqa: Multi-choice question answering benchmark for korean healthcare professional licensing examinations.](#)
- Keonmo Lee and Allganize Inc. LLM Team. 2024. Ifeval-ko: Korean instruction-following benchmark. <https://huggingface.co/datasets/allganize/IFEval-Ko>. A Korean translation of Google's IFEval dataset for evaluating instruction-following capabilities of LLMs using the lm-evaluation-harness framework.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.
- Chuang Liu, Linhao Yu, Jiaxuan Li, Renren Jin, Yufei Huang, Ling Shi, Junhui Zhang, Xinmeng Ji, Tingting Cui, Tao Liu, et al. 2024. Openeval: Benchmarking chinese llms across capability, alignment and safety. *arXiv preprint arXiv:2403.12316*.
- Meta AI. 2024. Llama 3.2 3b instruct. <https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct>. Version release date: 25 Sep 2024.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#).
- LG Research, Soyoung An, Kyunghoon Bae, Eunbi Choi, Kibong Choi, Stanley Jungkyu Choi, Seokhee Hong, Junwon Hwang, Hyojin Jeon, Gerrard Jeongwon Jo, et al. 2024. Exaone 3.5: Series of large language models for real-world use cases. *arXiv preprint arXiv:2412.04862*.
- Jaehyung Seo and Heuseok Lim. 2025. [K-HALU: Multiple answer korean hallucination benchmark for large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Hyopil Shin, Sangah Lee, Dongjun Jang, Wooseok Song, Jaeyoon Kim, Chaeyoung Oh, Hyemi Jo, Youngchae Ahn, Sihyun Oh, Hyohyeong Chang, Sunkyoung Kim, and Jinsik Lee. 2025. [Kobalt: Korean benchmark for advanced linguistic tasks.](#)
- Guijin Son, Hyunwoo Ko, and Dasol Choi. 2025. Multi-step reasoning in korean and the emergent mirage. *arXiv preprint arXiv:2501.05712*.
- Guijin Son, Hyunwoo Ko, Hoyoung Lee, Yewon Kim, and Seunghyeok Hong. 2024a. Llm-as-a-judge & reward model: What they can and cannot do. *arXiv preprint arXiv:2409.11239*.
- Guijin Son, Hanwool Lee, Sungdong Kim, Seung-gone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024b. Kmmlu: Measuring massive multitask language understanding in korean. *arXiv preprint arXiv:2402.11548*.
- Guijin Son, Hanwool Lee, Suwan Kim, Huiseo Kim, Jaecheol Lee, Je Won Yeom, Jihyu Jung, Jung Woo Kim, and Songseong Kim. 2023. Hae-rae bench: Evaluation of korean knowledge in language models. *arXiv preprint arXiv:2309.02706*.
- FlagEval Team. 2025. [Flageval: An open-source evaluation toolkit for large models](#). Accessed: March 10, 2025.
- 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](#).
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#).