

HanjaBridge: Resolving Semantic Ambiguity in Korean LLMs via Hanja-Augmented Pre-Training

Seungho Choi*, Sihyun Park*, Minsang Kim*, Chansol Park*, Bongsu Kim*

WisenuT

{csh1019, psh0430, minsang, chansol53, usgnob}@wisenuT.co.kr

Abstract

Large language models (LLMs) often show poor performance in low-resource languages like Korean, partly due to unique linguistic challenges such as homophonous Sino-Korean words that are indistinguishable in Hangul script. To address this semantic ambiguity, we propose HanjaBridge, a novel meaning-injection technique integrated into a continual pre-training (CPT) framework. Instead of deterministically mapping a word to a single Hanja (Chinese character), HanjaBridge presents the model with all possible Hanja candidates for a given homograph, encouraging the model to learn contextual disambiguation. This process is paired with token-level knowledge distillation to prevent catastrophic forgetting. Experimental results show that HanjaBridge significantly improves Korean language understanding, achieving a 21% relative improvement on the KoBALT benchmark. Notably, by reinforcing semantic alignment between Korean and Chinese through shared Hanja, we observe a strong positive cross-lingual transfer. Furthermore, these gains persist even when Hanja augmentation is omitted at inference time, ensuring practical efficiency with no additional run-time cost.

Code — <https://github.com/oh-gnues-iohc/HanjaBridge>

1 Introduction

Multilingual LLMs achieve impressive performance across various languages through joint training, (Moosa, Akhter, and Habib 2022) but these gains are not distributed equally among all languages (Ogueji, Zhu, and Lin 2021). In particular, languages with limited training corpora like Korean suffer a pronounced performance gap. During multilingual training, Korean often constitutes less than 0.1% of the data in many open-source LLMs, leading to notably lower performance compared to high-resource languages. Dedicated Korean-centric models (Kim et al. 2021) have been developed to mitigate this gap, but they still do not fully resolve Korean’s unique linguistic challenges.

Approximately 57% of Korean vocabulary is Sino-Korean (borrowed Chinese-origin words). (Park and Zhao 2019) These words share semantic and morphological roots with Chinese. For example, “연구” (Korean for “research”) shares

English: I started my **research** today.

Korean: 나는 오늘부터 **연구**를 시작했다.

Chinese: 我从今天开始**研究**。

English: I did some market **research** before launching the product.

Korean: 나는 제품 출시 전에 시장**조사**를 했다.

Chinese: 我在产品上市前做了市场**调查**。

Figure 1: Context-dependent Sino-Korean translations of the English word *research*

the same characters as Chinese “研究” (yánjiū), and in both languages one can form perfectly equivalent sentences like “나는 오늘 **연구**를 시작했다” (“I started my **research** today”) and “我今天**始**研究” This phenomenon is possible because Korean and Chinese share a common logographic representation (Hanja), implying that cross-lingual semantic alignment can be relatively straightforward at the word level when using these characters. In contrast, English, which uses an alphabetic script, often requires context to interpret the correct meaning of a word that has multiple senses.

Fig. 1. Cross-lingual variations in the translation of the English word “research.” The same English word is translated as “연구(研究)” in a scholarly context, whereas it becomes “조사(调查)” in a market-related context. This illustrates that although Chinese and Korean align semantically and morphologically via shared Sino-origin characters, English requires context-sensitive interpretation, making one-to-one mapping difficult.

However, modern written Korean uses the phonetic Hangul alphabet, which does not distinguish between different Hanja that share the same pronunciation. Thus, many distinct Sino-Korean words have identical Hangul spellings, creating a high degree of polysemy and homophony in the language. (Park and Zhao 2019) In fact, about 35% of Hanja entries in the Standard Korean Dictionary are homophones. For instance, the word “의사” can correspond to several different Hanja—“醫師” (doctor), “意思” (intention), “義士” (patriot), or “議事” (deliberation)—yet all are written “의사” in Hangul, making them virtually indistinguishable in text. This structural limitation of Hangul poses a major challenge for LLMs in accurately capturing word meanings and interpreting context. The model may struggle with semantic disambiguation since multiple meanings collapse into the same surface form.

*These authors contributed equally.

To address these issues, we propose HanjaBridge, a semantic augmentation method that leverages the logographic nature of Chinese characters to provide explicit meaning cues during training. By appending Hanja to Korean text and applying token-level knowledge distillation, our approach encourages the model to disambiguate meanings based on context while preserving its original multilingual capabilities.

We show that this method leads the model to increasingly focus attention on contextually appropriate Hanja candidates as training progresses (RQ1), and that it improves Korean semantic understanding without relying on Hanja at inference or degrading performance in other languages (RQ2). This confirms both the effectiveness and practicality of HanjaBridge for enhancing Korean LLMs.

2 Related Work

2.1 Continual Pre-Training for Low-Resource Languages

Large multilingual pre-trained language models (PLMs) exhibit strong performance across many languages, yet their effectiveness on low-resource languages (LRLs) remains limited (Liu, Winata, and Fung 2021; Guo et al. 2024). To address this imbalance, continual pre-training (CPT) has emerged as a core strategy, which further trains a multilingual model on large-scale monolingual corpora to enhance language-specific representation (Fujinuma, Boyd-Graber, and Kann 2022). CPT has been successfully applied to Korean, Japanese, and Southeast Asian languages (Vo et al. 2024; Fujii et al. 2024; Dou et al. 2024).

Recent studies extend beyond naive data continuation by improving training efficiency and stability. Parameter-efficient fine-tuning (PEFT) methods such as LoRA have been combined with CPT to reduce computational cost (Nag et al. 2025). Other approaches optimize the training process via data reweighting (Luo et al. 2024) or instruction-based CPT (Chen and Lee 2024). These advances indicate that CPT has evolved into a systematic framework for language specialization.

2.2 Tokenizer Vocabulary Expansion

A major obstacle to effective CPT is the vocabulary mismatch of multilingual tokenizers, which are primarily optimized for high-resource languages. In morphologically rich languages such as Korean, this leads to severe over-tokenization, where a single semantic unit is fragmented into multiple subword tokens (Lee et al. 2024; Kim et al. 2024). This semantic fragmentation increases sequence length and weakens the model’s ability to learn coherent word meaning.

Vocabulary expansion has been widely adopted to alleviate this issue by adding frequent characters or morphemes of the target language to the tokenizer. This allows important semantic units to be represented as single tokens, reducing token length and improving semantic consistency. Table 1 illustrates typical over-tokenization errors observed in Korean.

Original Word	Segmented Subwords
<p>골칫거리 (Nuisance / Headache)</p>	<p>골칫 + 거리 (Problem) + (Street / Distance)</p>
<p>눈치채다 (Notice / Perceive)</p>	<p>눈 + 치 + 채다 (Eye / Snow) + (Null) + (Snatch)</p>
<p>입덧 (morning sickness)</p>	<p>입 + 덧 (Mouth) + (Layer / Over)</p>

Table 1: Examples of over-tokenization in Korean causing meaning distortion.

2.3 Knowledge Distillation for Multilingual and Domain Adaptation

Although CPT and vocabulary expansion improve target-language performance, they often induce *catastrophic forgetting*, degrading performance in other languages due to the stability–plasticity dilemma. Knowledge distillation (KD) has proven to be an effective remedy in this setting (Xu et al. 2024). In KD, the original multilingual model serves as a teacher, and the continually trained model is optimized to match its output logits (Hinton, Vinyals, and Dean 2015).

Recent studies confirm that KD reliably preserves multilingual knowledge during low-resource language adaptation (Raju et al. 2025). Although model-merging approaches have also been proposed (Alexandrov et al. 2024; Tao et al. 2024), KD remains the most widely validated solution when explicit knowledge retention is required.

2.4 Resolving Semantic Ambiguity via External Signals

While CPT, vocabulary expansion, and KD form a strong foundation for language adaptation, they cannot fully resolve *inherent semantic ambiguity*. Korean contains numerous Sino-Korean homophones, where a single Hangul form maps to multiple meanings (e.g., “의사”: doctor, intention, patriot), making word-sense disambiguation particularly challenging (Navigli 2009).

To resolve such ambiguity, prior work has leveraged explicit *external semantic signals*. In the CJK context, converting Hangul to Hanja has consistently improved machine translation quality across Korean–Chinese, Korean–Japanese, and Vietnamese–Chinese translation tasks (Yoo, Kim, and Lee 2019; Kim, Hirasawa, and Komachi 2020; Li, Sha, and Shi 2020). These results suggest that exposing etymological and character-level semantics provides an effective mechanism for disambiguation and cross-lingual knowledge transfer.

3 Method

To enhance the semantic expressiveness of a Korean LLM, we propose HanjaBridge, whose overall workflow is illustrated in Fig. 2. Our approach consists of three key steps: (i) We constructed a joint Korean–Chinese dictionary of Sino-Korean words and added only Korean words that have corresponding Chinese characters to the multilingual tokenizer. (ii) These Hanja-annotated Korean words are appended in-line to the sentence so that the model must resolve the correct meaning

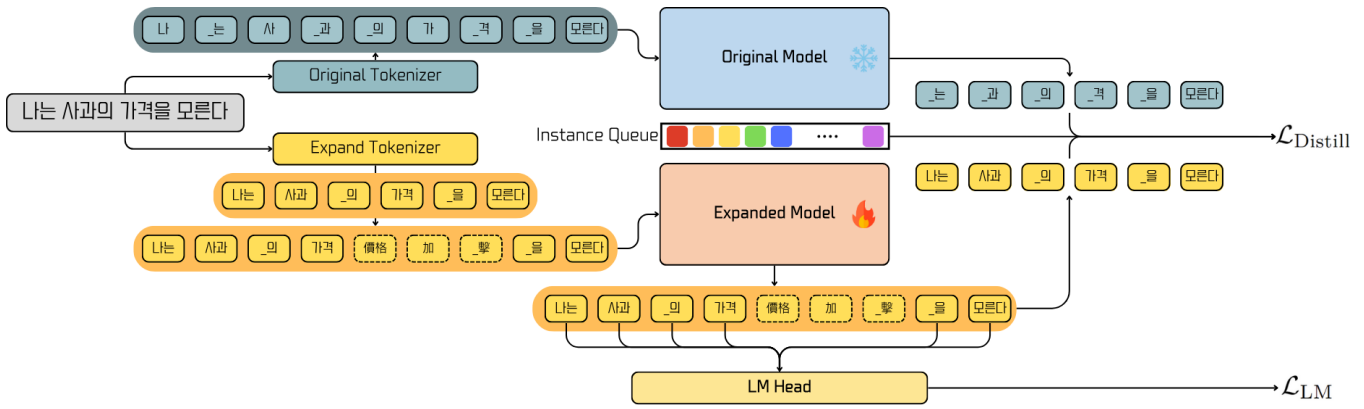


Figure 2: End-to-end workflow of our method.

from context. (iii) We apply token-wise knowledge distillation, forcing the student model to mimic the teacher’s output distributions and hidden representations, thereby boosting Korean performance without erasing existing multilingual knowledge.

3.1 HanjaBridge: A Hanja-Augmented Semantic Injection Method

HanjaBridge creates a single semantic slot by concatenating all candidate Hanja forms immediately after each ambiguous Hangul token. For example, the sentence “나는 사과의 가격을 모른다” (“I don’t know the price of the apple”) is internally expanded to “나는 사과의 가격價格 加 擊을 모른다,” where 價格 (‘price’) and 加擊 (‘hit’) are plausible Hanja spellings for 가격. During training, the model learns to select the contextually appropriate Hanja inside this slot, thereby disambiguating the original Korean word.

Because multiple Hanja candidates are presented simultaneously, the model cannot merely memorize a single correct answer; instead, it must leverage *context* to decide which Hanja yields the highest-probability distribution. This not only trains the model’s innate disambiguation ability but also injects the rich semantic and etymological information encoded in high-resource Chinese directly into Korean tokens. Consequently, the model (1) interprets homophonous Korean words correctly in context, and (2) exploits Chinese semantic knowledge when processing Korean, leading to broad gains in expressiveness, reasoning accuracy, and knowledge coverage.

The model’s attention mechanism is also adjusted to match these expanded inputs, as shown in Fig. 3. Within each expansion group, the Korean token can freely attend to its Hanja candidates, but *inter-candidate* attention is blocked. From a Korean perspective, these Hanja share the same surface form, yet from a Chinese perspective they differ semantically; allowing candidates to attend to one another would blur their meanings. By restricting the attention pathway to “Korean token \leftrightarrow each Hanja candidate,” we ensure that every candidate retains its independent meaning while still informing the Korean token’s representation.

During training, Hanja tokens participate *only* as semantic

	나는	사과	의	가격	價格	加	擊	을	모른다
나는	1	0	0	0	0	0	0	0	0
사과	1	1	0	0	0	0	0	0	0
의	1	1	1	0	0	0	0	0	0
가격	1	1	1	1	1	1	1	0	0
價格	1	1	1	1	1	0	0	0	0
加	1	1	1	1	0	1	1	0	0
擊	1	1	1	1	0	1	1	0	0
을	1	1	1	1	1	1	1	1	0
모른다	1	1	1	1	1	1	1	1	1

Figure 3: This is an illustration of the attention mask of HanjaBridge applied to the example sentence. Blue cells indicate the original baseline causal mask, while yellow cells highlight the additional connections introduced by HanjaBridge.

hints; the language-modeling loss is computed exclusively on the original Korean tokens. Let the input length be L and assign a binary mask $m_t \in \{0, 1\}$ to each position t . Positions with $m_t = 1$ are original tokens; $m_t = 0$ are Hanja expansions. Define the set of original positions as $O = \{t \mid m_t = 1\}$. Only hidden states h_t ($t \in O$) are fed into the LM head:

$$\mathbf{Z}_t = \mathbf{W}h_t + \mathbf{b}, \quad P(y_t \mid \mathbf{x}_{<t}) = \text{softmax}(\mathbf{Z}_t). \quad (1)$$

$$\mathcal{L}_{\text{LM}} = - \sum_{t \in O} \log P(y_t \mid \mathbf{x}_{<t}). \quad (2)$$

Thus, Hanja tokens act only as latent cues via attention, while prediction and parameter updates focus on the original sequence.

3.2 Token-wise Knowledge Distillation

To preserve multilingual competence, we adopt token-level knowledge distillation (Fig. 4). The pretrained source model serves as a *frozen teacher*. The student is parameter-initialized from the teacher but unfreezes only selected layers—e.g. embeddings and a subset of Transformer blocks—enabling efficient training and mitigating catastrophic forgetting.




Figure 4: The training strategy for the Teacher-Student models. All parameters of the Teacher model are frozen and used solely as a knowledge source. The Student model is trained only on select layers and the newly expanded embeddings.

Within this teacher–student setup, we maintain an *instance queue* that stores the teacher model’s token-level hidden vectors. The student is then trained to replicate these vectors at the feature level. Because the two models use different tokenizers, their sequence lengths can mismatch and the teacher may split a single word into multiple sub-words. We resolve this misalignment via *offset mapping*: every original token position in the student’s input is put into one-to-one correspondence with the teacher, and when a teacher token is fragmented into several sub-words, we take the hidden vector of the *last* sub-word as its representative. Hanja tokens serve only as auxiliary context and are therefore excluded from the distillation loss.

Our method distills knowledge by aligning *token-level hidden representations*, not sequence-level summaries or output logits. This token-wise distillation offers two advantages. First, every token supplies a learning signal, enabling dense supervision. Second, even a single sentence adds dozens to hundreds of tokens to the queue, so the number of positive and negative examples grows explosively; this lets the student estimate the teacher’s feature space with far greater fidelity. As a result, the student learns to imitate not just individual vectors but the *structure* of the feature space itself, preserving contextual meaning and relational information more effectively.

To steer the student’s feature learning effectively, we adopt a *contrastive, queue-based distillation* strategy. (Fang et al. 2021) We maintain a fixed-length instance queue D : after each mini-batch, the teacher model’s token-level hidden vectors are enqueued, and the oldest entries are dequeued so that the queue size stays constant. During training, for every original token i we align the teacher vector z_i^T with the corresponding student vector z_i^S . We then form the augmented set $D^+ = D \cup \{z_i^T\}$ —that is, the current teacher vector temporarily pushed onto the queue—and compute the similarity distributions $p_T(i)$ and $p_S(i)$ as defined in Eq. 3. When $z_i = z_i^T$ we obtain p_T ; when $z_i = z_i^S$ we obtain p_S . The student consults *only* teacher vectors, never its own past outputs, and minimizes the cross-entropy between $p_T(i)$ and $p_S(i)$. This procedure supplies a rich *multi-point alignment* signal: the student learns not merely to copy each teacher vector in isolation, but also to reproduce its *relative relationships* to every other token in the queue, enabling more comprehensive knowledge transfer.

$$p(i, j) = \frac{\exp((z_i \cdot d_j)/\tau)}{\sum_{d \in D^+} \exp((z_i \cdot d)/\tau)}. \quad (3)$$

The overall training procedure jointly optimizes two losses: (1) the standard language-modeling loss for the student, L_{LM} , and (2) the token-level distillation loss L_{KD} defined in Eq. 4. Both losses are computed *only* at positions corresponding to the original Korean tokens; Hanja tokens are excluded from training.

$$L_{KD} = -\frac{1}{N} \sum_{i=1}^N \sum_{d_j \in D^+} p_T(i, j) \log p_S(i, j). \quad (4)$$

The final objective is the weighted sum of the two losses, as shown in Eq. 5:

$$L_{total} = L_{LM} + \lambda L_{KD}. \quad (5)$$

Here, λ is a hyper-parameter that balances the gradient magnitudes of the two terms. This parallel optimization enables the student to internalize not only raw language-modeling ability but also the teacher’s representational structure. Empirically, the proposed distillation scheme yields larger performance gains than simple logit matching, demonstrating that the combination of Hanja-based auxiliary information with dense feature supervision is an effective vehicle for knowledge transfer.

4 Experiments

4.1 Experimental Setup

Hardware. All training was conducted on a system with $4 \times$ NVIDIA H200 80GB GPUs, using mixed precision (BF16) training for efficiency.

Base model. We build on the open-source **Qwen 2.5-3B** model as our base. We apply our proposed HanjaBridge augmentation, tokenizer expansion, and token-level distillation to this model as described in Sec. 3.

Continual pre-training data. We use only **Korean data** for continual pre-training. (No additional English or Chinese data was used in this phase.) The corpus is drawn from the FineWeb-Edu dataset (Penedo et al. 2024), a high-quality Korean text collection.

Hyper Parameters. We set the sequence length to 65,536 tokens. For knowledge distillation, we use a temperature of $\tau_T = 0.01$ for the teacher and $\tau_S = 0.2$ for the student, and a distillation loss weight of $\lambda = 0.1$.

Benchmarks.

- **Korean:** KoBEST (Jang et al. 2022), KoBALT (Shin et al. 2025)
- **English:** BoolQ (Clark et al. 2019), COPA (Roemmele, Bejan, and Gordon 2011), Hellaswag (Zellers et al. 2019), WiC (Pilehvar and Camacho-Collados 2019)

KoBEST is an evaluation suite of 5 Korean downstream tasks requiring linguistic knowledge, constructed with high-quality human-generated data. KoBALT is a comprehensive

Model	Korean (KoBALT - Hard)						Korean (KoBEST - General)						English				
	syntax	semantics	pragmatics	morphology	phonetics	avg	boolq	copa	hellaswag	sentineg	wic	avg	boolq	copa	hellaswag	wic	avg
Qwen2.5-3B	0.1067	0.2047	0.0988	0.0714	0.1129	0.1189	0.5413	0.6340	0.4600	0.4962	0.4229	0.5108	0.7722	0.8500	0.5505	0.6176	0.6976
Full CPT	0.1000	0.1907	0.0741	0.0476	0.1452	0.1115	0.7236	0.6450	0.4460	0.6080	0.5887	0.6023	0.7471	0.8200	0.5020	0.5313	0.6501
k=0 (distill only)	0.0900	0.1581	0.0741	0.0952	0.0968	0.1028	0.7201	0.6470	0.4560	0.5668	0.4550	0.5690	0.7731	0.8300	0.5089	0.5893	0.6753
k=2	0.1033	0.2000	0.0998	0.0238	0.0645	0.0983	0.7137	0.6520	0.4480	0.5743	0.6085	0.5993	0.7865	0.8200	0.5068	0.5846	0.6745
k=4	0.1033	0.2186	0.1358	0.0238	0.1129	0.1189	0.6830	0.6600	0.4500	0.6650	0.6064	0.6129	0.7662	0.8200	0.5087	0.5172	0.6531
k=8	0.1067	0.2279	0.1358	0.1190	0.1290	0.1437	0.7244	0.6560	0.4500	0.7380	0.6087	0.6354	0.7788	0.8300	0.5060	0.6019	0.6792
k=16	0.1067	0.2140	0.1235	0.0714	0.1452	0.1322	0.6980	0.6700	0.4480	0.6952	0.6086	0.6240	0.7725	0.8300	0.5059	0.5846	0.6733

Table 2: Comparison of model performance on the main benchmarks. *KoBALT* is a suite of difficult tasks that require deep linguistic understanding—syntax, semantics, and so on—whereas *KoBEST* measures general natural-language-understanding (NLU) ability. The variable k denotes the maximum number of Hanja candidates supplied during HanjaBridge training. Boldface marks the highest average (*avg*) score within each benchmark group.

Korean language understanding benchmark covering five linguistic domains (syntax, semantics, pragmatics, phonology, morphology). For English, we use widely adopted benchmarks that test commonsense reasoning and contextual understanding: BoolQ, COPA, HellaSwag, and WiC. All tasks are evaluated in zero-shot mode (no fine-tuning), and accuracy is used as the evaluation metric for all tasks.

4.2 Main Results: Korean vs English

We evaluate our proposed HanjaBridge method on both Korean and English benchmarks to assess overall performance. We compare the optimal configuration (our model with $k = 8$ Hanja candidates) against the original pre-trained model and two baselines: a standard continual pre-training without our semantic augmentation (Full CPT), and a CPT with knowledge distillation but no Hanja augmentation ($k=0$). Table 2 summarizes the results.

Korean Language Understanding : HanjaBridge substantially improves the model’s Korean understanding abilities, especially on tasks requiring nuanced semantic disambiguation. On the challenging KoBALT-Hard benchmark that evaluates deep linguistic knowledge, our $k=8$ model achieves an average score of 0.1437, outperforming the baseline (0.1189) by about 21% relative improvement. This demonstrates that providing Hanja candidates helps the model grasp complex syntactic and semantic structures more effectively. Similarly, on the more general natural language understanding suite KoBEST-General, our model records the highest average score of 0.6354, outperforming all other models. These results validate our hypothesis that explicitly injecting semantic candidates via HanjaBridge is a highly effective strategy for enhancing Korean-specialized capabilities.

Preventing Catastrophic Forgetting and Cross-Language Transfer One of the key issues in continuous learning for language specialization is catastrophic forgetting, where models lose their ability to understand other languages. Without separate mitigation strategies, the Full CPT model shows a significant decline in English performance from an average of 0.6976 to 0.6501, clearly demonstrating the problem of forgetting. In contrast, our methodology successfully addresses this issue. The $k=8$ model not only demonstrates outstanding performance in Korean but also effectively preserves English performance, achieving an average score of **0.6792**. This score is significantly higher than the bench-

mark of the Full CPT model and even slightly exceeds the performance of the $k=0$ (distill only) model. This demonstrates that meaning enhancement through Hanja does not come at the expense of English proficiency.




Figure 5: Cumulative attention heat-maps (36 heads \times 12 layers) for two sentences containing the ambiguous Korean word “의사.” Green highlights mark the context-appropriate Hanja (醫師 ‘doctor’; 意思 ‘intention’). (A) HanjaBridge focuses strongly on the correct character, whereas the (B) Pretrained-Only model distributes attention almost uniformly across all candidates.

Effect of the number of Chinese character candidates (k) Through experiments, we also discovered a trend regarding the number of Chinese character candidates (k). The Korean benchmark performance improved overall as k increased from 2 to 8. However, this trend did not continue indefinitely, and the $k=16$ model showed a slight performance decline

k (Number of Candidates)	Training Steps					
	30k	60k	90k	120k	150k	180k
2	0.573	0.581	0.597	0.611	0.631	0.638
3	0.391	0.413	0.427	0.446	0.453	0.477
4	0.321	0.338	0.347	0.362	0.375	0.404
5	0.267	0.286	0.324	0.349	0.360	0.388
6	0.108	0.146	0.209	0.267	0.323	0.361
7	0.055	0.108	0.125	0.186	0.234	0.285
8	0.102	0.148	0.177	0.239	0.277	0.308

Table 3: RQ1: Accuracy of selecting the correct Hanja candidate across training steps. k indicates the number of Hanja candidates, and values represent the proportion where the highest attention was assigned to the correct candidate.

compared to the $k=8$ model. This suggests that while providing candidate options is beneficial, an excessive number of candidates can introduce noise and complicate the model’s semantic discrimination process.

4.3 RQ1 – Attention-focus Analysis

In RQ1, we analyzed whether the proposed Chinese character pairing technique was correctly reflected in the model’s internal attention. To verify this, we analyzed the self-attention weight patterns of the model trained after Chinese character pairing. Specifically, we observed whether the model’s attention was focused on the correct candidate among the multiple Chinese character candidates input together with the Korean word.

However, the raw attention map of the Transformer layer tends to be flattened, making direct interpretation difficult. This is because, as the layers increase, each token’s embedding becomes dependent on the overall context, making it difficult to understand how information is propagated using simple weights alone. Therefore, instead of simple attention probabilities, we utilized propagation-based contribution analysis methods such as Attention Rollout. This method calculates the cumulative effect of attention across multiple layers from input to output, ultimately tracking the extent to which a specific input token contributed to the formation of the target output embedding.

Fig. 5 shows the results of visualizing the Attention Rollout of the HanjaBridge model and the original (Pretrained-Only) model after accumulating 36 layers, using two Korean sentences containing homophonic Chinese characters such as “의사 (physician/intention/deliberation).” All Hanja candidates following the Korean tokens were input, and the green characters indicate the correct Hanja for each context. The model trained using Hanja annotation focuses attention more strongly on the correct Hanja characters (醫師, 意思) in both sentences, while the original model distributes attention almost evenly across all candidates. This visually supports the claim that expanded Hanja tokens actually enhance meaning representation.

Table 3 shows the results of quantitative evaluations performed on tens of thousands of sentences using experiments such as Fig. 5, demonstrating how much attention the model focuses on the correct character candidates (correct character tokens) during the training steps (30k ~ 180k). The values

다음 문장에서 괄호 [] 안의 한국어 단어에 대응하는 올바른 한자 표기를 보기 중에서 선택 하시오.

문장: [의사]의 진찰을 받다.

보기:

A. 醫師 B. 意思 C. 議事

정답:

Choose the correct Chinese character representation for the Korean word inside the brackets [] in the following sentence.

Sentence: I received a check-up from the [doctor].

A. doctor B. intention C. discussion

Answer:

Figure 6: Examples of evaluation prompts ($k = 3$).

in the table represent the “accuracy” (the percentage of correct candidates receiving the highest attention) within each difficulty level ($k =$ number of candidates) interval.

The experimental results show that accuracy steadily increases in all cases as learning progresses. For example, in the easy interval ($k = 2$), accuracy increased from 0.573 to 0.638, and in the difficult interval ($k = 8$), accuracy increased from 0.102 to 0.308, confirming that the model gradually utilizes Hanja expansion information more actively. Additionally, differences in convergence speed depending on difficulty are observed. In the easy case, the improvement in performance slows down after a certain amount of learning, while in the difficult case, the upward trend continues even at the end of learning. This suggests that convergence should be considered for each difficulty level when designing the learning schedule.

In summary, the RQ1 experiment strongly supports the hypothesis that the proposed method actually utilizes Hanja information at the attention level and that the tendency to focus on correct Hanja characters is reinforced as learning progresses. In particular, the consistent improvement in high-difficulty samples supports the core claim that Hanja-Bridge and token-level distillation effectively enhance the semantic discrimination ability of Korean LLMs.

4.4 RQ2 – Hanja Multiple-choice Probe

Evaluation Setup We measured whether the model could select the correct Chinese character notation corresponding to *homophonic Korean words* based on context. The prompt format is as shown below, and the model is tasked with solving a multiple-choice question by selecting the correct Chinese character from k options (candidate Chinese characters) in a zero-shot manner.

The model score is calculated using the log-likelihood \log_p for each option. To eliminate **bias due to option order**, k prompts were generated so that the correct candidate appeared once at each position, and the average log-likelihood for each option was used as the final score.

Comparison Settings

- **HanjaBridge (w/ HB-inf.)** : HanjaBridge applied to both learning and inference
- **HanjaBridge (without HB-inf.)** : HanjaBridge applied during training, inference uses Korean only

Model	Accuracy	Total Tokens	Avg Tokens / Sample
HanjaBridge (w HB-inf.)	0.575	961,982	96.19
HanjaBridge (w/o HB-inf.)	0.568	922,384	92.23
Full CPT	0.389	922,384	92.23
Pretrain	0.503	1,030,151	103.01

Table 4: Accuracy comparison based on Chinese character merging method and inference conditions

- **Full CPT** : Full-parameter CPT performed using Korean corpus only
- **Pretrain** : Original pre-trained model Qwen2.5-3B

Analysis and Implications (i) In Table 4, the proposed model outperforms pretraining (0.503) by 6.5 percentage points and Full CPT (0.389) by 17.9 percentage points even when only using Chinese character combinations in the training stage (0.568). In other words, once Hanja merging is used to inject semantic discrimination information during training, it maintains its advantage during inference without requiring additional tokens. This means that both **inference efficiency**—no increase in token length—and **model versatility**—can be deployed as-is on existing APIs and servers—are secured simultaneously.

(ii) Pre-trained models already possess a certain level of Korean–Chinese character meaning connection, but *Chinese character merging CPT* further **concretizes and reinforces this at the lexical level**, raising it to the highest value of 0.575. This provides quantitative evidence that the tendency observed in RQ1, where attention focuses on the correct Chinese character, also translates into actual task performance.

(iii) The setting that uses merged tokens for both learning and inference is slightly higher (+0.7 percentage points), but the difference is not significant, so in actual service, the “merge only during learning” mode is a reasonable choice in terms of cost-performance balance.

In summary, RQ2 demonstrates that Hanja merging CPT is a practical method that enhances Korean semantic resolution while allowing the existing tokenizer and system to be used as-is during the inference stage.

5 Conclusion

In this study, we proposed **HanjaBridge**, a continual pre-training framework that addresses semantic ambiguity in Korean caused by Sino-Korean homophones. Instead of injecting a single fixed Hanja, HanjaBridge provides multiple candidates in parallel, enabling contextual disambiguation through learning. This process is combined with token-level knowledge distillation and vocabulary expansion to prevent catastrophic forgetting.

Experimental results demonstrate that HanjaBridge achieves significant improvements on Korean language understanding benchmarks, including a **21% relative improvement** on KoBALT. At the same time, knowledge distillation effectively preserves existing multilingual capabilities, enabling positive cross-lingual transfer. Attention analysis further confirms that the model progressively attends to contextually correct Hanja candidates during training.

A key practical advantage of HanjaBridge is that these gains are achieved without additional inference-time cost. Even when Hanja is used only during training and removed at inference, performance remains stable, allowing seamless deployment.

While HanjaBridge relies on the quality of the Hanja-Hangul dictionary and may be limited for neologisms or domain-specific terms, future work will explore the impact of dictionary scale and quality to guide cost-effective domain adaptation.

Author Contributions

Author	Contributions
Seunggho Choi	His contributions to this research include proposing and implementing the detailed methodology, conducting theoretical analysis of the core algorithms, managing the entire experimental pipeline, and writing the methodology section. He also reviewed the entire manuscript.
Sihyun Park	His contributions include designing training strategies, proposing model improvement directions, and conducting preliminary experiments for hyperparameter tuning. He drafted the introduction section and reviewed the logical structure and overall flow of the manuscript.
Minsang Kim	His contributions include designing auxiliary experiments, deriving insights, conducting experiments, and organizing results. He drafted the experiment section and reviewed and edited the entire manuscript.
Chansol Park	His contributions include designing and iteratively refining the evaluation methodology, constructing dictionaries, designing evaluation datasets, and computing baseline performance. He drafted the related works section and reviewed the full manuscript.
Bongsu Kim	His contributions include conceptualizing the framework and designing the overall structure, providing project-level guidance and supervision, and reviewing the consistency and completeness of all sections. He supervised the entire work and ensured coherence across the manuscript.

Table 5: Author Contributions

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2710017875, Development of multimodal data input-based search augmentation generation technology).

References

- Alexandrov, A.; Raychev, V.; Müller, M. N.; Zhang, C.; Vechev, M.; and Toutanova, K. 2024. Mitigating Catastrophic Forgetting in Language Transfer via Model Merging. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 17167–17186. Miami, Florida, USA: Association for Computational Linguistics.
- Chen, K.-M.; and Lee, H.-y. 2024. InstructionCP: A fast approach to transfer Large Language Models into target language. *arXiv preprint arXiv:2405.20175*.

- Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2924–2936. Minneapolis, Minnesota: Association for Computational Linguistics.
- Dou, L.; Liu, Q.; Zeng, G.; Guo, J.; Zhou, J.; Mao, X.; Jin, Z.; Lu, W.; and Lin, M. 2024. Sailor: Open Language Models for South-East Asia. In Hernandez Farias, D. I.; Hope, T.; and Li, M., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 424–435. Miami, Florida, USA: Association for Computational Linguistics.
- Fang, Z.; Wang, J.; Wang, L.; Zhang, L.; Yang, Y.; and Liu, Z. 2021. {SEED}: Self-supervised Distillation For Visual Representation. In *International Conference on Learning Representations*.
- Fujii, K.; Nakamura, T.; Loem, M.; Iida, H.; Ohi, M.; Hattori, K.; Shota, H.; Mizuki, S.; Yokota, R.; and Okazaki, N. 2024. Continual Pre-Training for Cross-Lingual LLM Adaptation: Enhancing Japanese Language Capabilities. In *First Conference on Language Modeling*.
- Fujinuma, Y.; Boyd-Graber, J.; and Kann, K. 2022. Match the Script, Adapt if Multilingual: Analyzing the Effect of Multilingual Pretraining on Cross-lingual Transferability. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1500–1512. Dublin, Ireland: Association for Computational Linguistics.
- Guo, P.; Ren, Y.; Hu, Y.; Li, Y.; Zhang, J.; Zhang, X.; and Huang, H. 2024. Teaching Large Language Models to Translate on Low-resource Languages with Textbook Prompting. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 15685–15697. Torino, Italia: ELRA and ICCL.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Jang, M.; Kim, D.; Kwon, D. S.; and Davis, E. 2022. KoBEST: Korean Balanced Evaluation of Significant Tasks. In Calzolari, N.; Huang, C.-R.; Kim, H.; Pustejovsky, J.; Wanner, L.; Choi, K.-S.; Ryu, P.-M.; Chen, H.-H.; Donatelli, L.; Ji, H.; Kurohashi, S.; Paggio, P.; Xue, N.; Kim, S.; Hahm, Y.; He, Z.; Lee, T. K.; Santus, E.; Bond, F.; and Na, S.-H., eds., *Proceedings of the 29th International Conference on Computational Linguistics*, 3697–3708. Gyeongju, Republic of Korea: International Committee on Computational Linguistics.
- Kim, B.; Kim, H.; Lee, S.-W.; Lee, G.; Kwak, D.; Dong Hyeon, J.; Park, S.; Kim, S.; Kim, S.; Seo, D.; Lee, H.; Jeong, M.; Lee, S.; Kim, M.; Ko, S. H.; Kim, S.; Park, T.; Kim, J.; Kang, S.; Ryu, N.-H.; Yoo, K. M.; Chang, M.; Suh, S.; In, S.; Park, J.; Kim, K.; Kim, H.; Jeong, J.; Yeo, Y. G.; Ham, D.; Park, D.; Lee, M. Y.; Kang, J.; Kang, I.; Ha, J.-W.; Park, W.; and Sung, N. 2021. What Changes Can Large-scale Language Models Bring? Intensive Study on HyperCLOVA: Billions-scale Korean Generative Pretrained Transformers. In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3405–3424. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Kim, H.; Hirasawa, T.; and Komachi, M. 2020. Korean-to-Japanese Neural Machine Translation System using Hanja Information. In Nakazawa, T.; Nakayama, H.; Ding, C.; Dabre, R.; Kunchukuttan, A.; Pa, W. P.; Bojar, O.; Parida, S.; Goto, I.; Mino, H.; Manabe, H.; Sudoh, K.; Kurohashi, S.; and Bhattacharyya, P., eds., *Proceedings of the 7th Workshop on Asian Translation*, 127–134. Suzhou, China: Association for Computational Linguistics.
- Kim, S.; Park, J.; Kim, Y.; and Lee, S. 2024. KOMBO: Korean Character Representations Based on the Combination Rules of Subcharacters. In Ku, L.-W.; Martins, A.; and Sriku-mar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 5102–5119. Bangkok, Thailand: Association for Computational Linguistics.
- Lee, J.; Moon, H.; Lee, S.; Park, C.; Eo, S.; Ko, H.; Seo, J.; Lee, S.; and Lim, H. 2024. Length-aware Byte Pair Encoding for Mitigating Over-segmentation in Korean Machine Translation. In Ku, L.-W.; Martins, A.; and Sriku-mar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 2287–2303. Bangkok, Thailand: Association for Computational Linguistics.
- Li, H.; Sha, J.; and Shi, C. 2020. Revisiting back-translation for low-resource machine translation between Chinese and Vietnamese. *IEEE Access*, 8: 119931–119939.
- Liu, Z.; Winata, G. I.; and Fung, P. 2021. Continual Mixed-Language Pre-Training for Extremely Low-Resource Neural Machine Translation. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2706–2718. Online: Association for Computational Linguistics.
- Luo, Z.; Zhang, X.; Liu, X.; Li, H.; Gong, Y.; Qi, C.; and Cheng, P. 2024. Velocitune: A Velocity-based Dynamic Domain Reweighting Method for Continual Pre-training. *arXiv preprint arXiv:2411.14318*.
- Moosa, I. M.; Akhter, M. E.; and Habib, A. B. 2022. Transliteration: A Simple Technique For Improving Multilingual Language Modeling.
- Nag, A.; Chakrabarti, S.; Mukherjee, A.; and Ganguly, N. 2025. Efficient Continual Pre-training of LLMs for Low-resource Languages. In Chen, W.; Yang, Y.; Kachuee, M.; and Fu, X.-Y., eds., *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: Industry Track)*, 304–317. Albuquerque, New Mexico: Association for Computational Linguistics. ISBN 979-8-89176-194-0.
- Navigli, R. 2009. Word sense disambiguation: A survey. *ACM Comput. Surv.*, 41(2).

- Ogueji, K.; Zhu, Y.; and Lin, J. 2021. Small Data? No Problem! Exploring the Viability of Pretrained Multilingual Language Models for Low-resourced Languages. In Ataman, D.; Birch, A.; Conneau, A.; Firat, O.; Ruder, S.; and Sahin, G. G., eds., *Proceedings of the 1st Workshop on Multilingual Representation Learning*, 116–126. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Park, J.; and Zhao, H. 2019. Korean-to-chinese machine translation using chinese character as pivot clue. *arXiv preprint arXiv:1911.11008*.
- Penedo, G.; Kydlíček, H.; Lozhkov, A.; Mitchell, M.; Raffel, C. A.; Von Werra, L.; Wolf, T.; et al. 2024. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37: 30811–30849.
- Pilehvar, M. T.; and Camacho-Collados, J. 2019. WiC: the Word-in-Context Dataset for Evaluating Context-Sensitive Meaning Representations. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 1267–1273. Minneapolis, Minnesota: Association for Computational Linguistics.
- Raju, J. S.; Walia, J. S.; Raghav, S.; Marivate, V.; et al. 2025. AfroXLMR-Comet: Multilingual Knowledge Distillation with Attention Matching for Low-Resource languages. *arXiv preprint arXiv:2502.18020*.
- Roemmele, M.; Bejan, C. A.; and Gordon, A. S. 2011. Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning. In *AAAI spring symposium: logical formalizations of commonsense reasoning*, 90–95.
- Shin, H.; Lee, S.; Jang, D.; Song, W.; Kim, J.; Oh, C.; Jo, H.; Ahn, Y.; Oh, S.; Chang, H.; et al. 2025. KoBALT: Korean Benchmark For Advanced Linguistic Tasks. *arXiv preprint arXiv:2505.16125*.
- Tao, M.; Zhang, C.; Huang, Q.; Ma, T.; Huang, S.; Zhao, D.; and Feng, Y. 2024. Unlocking the Potential of Model Merging for Low-Resource Languages. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 8705–8720. Miami, Florida, USA: Association for Computational Linguistics.
- Vo, A.-D.; Jung, M.; Lee, W.; and Choi, D. 2024. Redwhale: An adapted korean llm through efficient continual pretraining. *arXiv preprint arXiv:2408.11294*.
- Xu, X.; Li, M.; Tao, C.; Shen, T.; Cheng, R.; Li, J.; Xu, C.; Tao, D.; and Zhou, T. 2024. A survey on knowledge distillation of large language models. *arXiv preprint arXiv:2402.13116*.
- Yoo, K. M.; Kim, T.; and Lee, S.-g. 2019. Don't Just Scratch the Surface: Enhancing Word Representations for Korean with Hanja. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3528–3533. Hong Kong, China: Association for Computational Linguistics.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4791–4800. Florence, Italy: Association for Computational Linguistics.