NMIXX vs. Domain-Adaptive Cross-Lingual Embeddings in Korean Financial QA

Assignee Research

June 3, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: How does the performance of NMIXX compare to domain-adaptive cross-lingual embeddings like LaBSE or XLM-R on financial QA benchmarks in Korean when evaluated using semantic similarity metrics beyond. 11 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: NMIXX: Domain-Adapted Neural Embeddings for Cross-Lingual eXploration of Finance. Research question: How does the performance of NMIXX compare to domain-adaptive cross-lingual embeddings like LaBSE or XLM-R on financial QA benchmarks in Korean when evaluated using semantic similarity metrics beyond cosine similarity?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.8/10.

3 Results

11 papers retrieved. 11 claims extracted; 0 independently verified. Quality review score: 3.8/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The training corpus consists of 18.8k high-confidence triplets , ,.	×	0.11
Hard negatives () are generated according to our four-axis financial semantic-shift taxonomy (temporal drift, perspecti	×	0.07
Positives () combine two complementary strategies: (1) In-domain paraphrase and (2) Exact translation.	×	0.07
We adopt a temperature-scaled triplet negative-log-likelihood loss.	×	0.02
Baseline Models are selected based on the intersection of those evaluated in the MTEB [19], FinMTEB [27], and KorFinMTEB	×	0.03
We fine-tune seven license-compatible embedding models that can be run on our four-A100 setup.	×	0.07
We measure Spearman's in four STS suites: FinSTS, KorFinSTS, STS, and KorSTS.	×	0.08
Training Configuration includes four NVIDIA A100 GPUs, AdamW optimizer with a fixed learning rate of 5 $\times$ 10-5, linear sc	×	0.04
Domain Adaptation Trade-offs show that domain-specific fine-tuning can erode performance in general-domain tasks.	×	0.10
Financial texts present linguistic and structural challenges such as domain-specific jargon, temporal semantic drift, an	×	0.07
Continual and supervised pre-training frameworks like FinGPT and BloombergGPT address domain specificity.	×	0.04

References

- http://arxiv.org/abs/1809.08928v1
- http://arxiv.org/abs/2507.09601v2
- http://arxiv.org/abs/2502.14620v1