One prompt is not enough: Instruction Sensitivity Undermines Embedding Model Evaluation

Yevhen Kostiuk*
Aarhus University
ykost@cas.au.dk

Kenneth Enevoldsen*
Aarhus University
kenneth.enevoldsen@cas.au.dk

Abstract

Instruction embedding models have become common among state-of-the-art models, however are evaluated using a single prompt per task. The single-point evaluation ignores a main problem of the instruction-based approach namely: sensitivity to the phrasing of the instruction. We present an empirical study of prompt sensitivity across 6 embedding models and 11 datasets. We show that reported scores misrepresent the distribution of scores over plausible prompts. The default prompt can both systematically understate or overstate performance. Furthermore, we show that the leaderboard ranking is not robust to prompt selection: a developer improve their rank by favorably selecting prompts, and under adversarial prompt selection any model can be promoted to first place. Our findings suggest that single-prompt evaluation is insufficient for instruction-tuned embedding models and that benchmarks should incorporate prompt robustness, either by evaluating over multiple prompts or by reporting sensitivity alongside point estimates.

The code can be found in our GitHub repository[1].

1 Introduction and Background

Embeddings are foundational to modern NLP applications, serving as the backbone for tasks such as Retrieval-Augmented Generation (RAG), semantic search, zero-shot clustering as well as newer tasks such as agentic memory (Lewis et al., 2020; Xu et al., 2026). Recently, instruction-tuned embedding models initially popularized by Su et al. (2023) have become common among top-performing embedding models such as Qwen or the e5-series (Zhang et al., 2025; Wang et al., 2024; Enevoldsen et al., 2025). Instructions (or prompts) enables models to capture task-specific nuances of the text, offering higher flexibility for various domains and objectives. However, while instructions allow for a high degree of user-flexibility they also allow instructions to be tuned to obtain better benchmark performance, thus posing a fundamental challenge of how to best measure, report and compare such models that do not oversell nor undersell their capabilities.

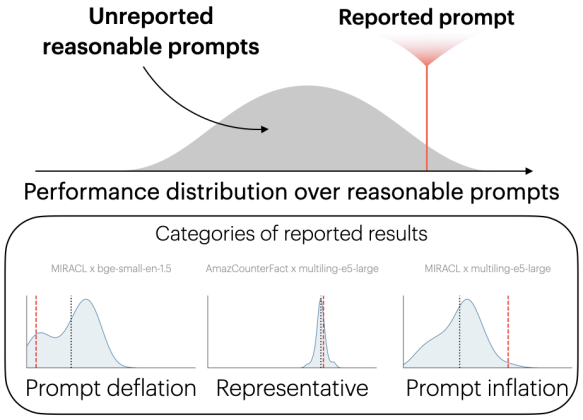


Figure 1: **Overview of our approach**, showing the currently unreported distribution of performances over reasonable prompts against the point estimate that is the reported results. Below that we show three possible option, either the reported result is representative of the distribution, it is lower than the (prompt deflation) or higher than (prompt inflation).

The most common benchmarks for evaluating the quality of embedding are implemented using the `beir` for information retrieval (Thakur et al., 2021; Kamalloo et al., 2024) and multitask `mteb` frameworks, which includes popular benchmarks such as MTEB, MMTEB (Muennighoff et al., 2022; Enevoldsen et al., 2025). `mteb` evaluate instruction-based models using only a single prompt either supplied by the task or the model contributor. This approach ignores a defining characteristic of these models: that they are sensitive to phrasing of the instruction. Relying on a single

*Equal contribution.

[1] https://github.com/centre-for-humanities-computing/one-prompt-is-not-enough

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 19408–19427
October 24-29, 2026 ©2026 Association for Computational Linguistics

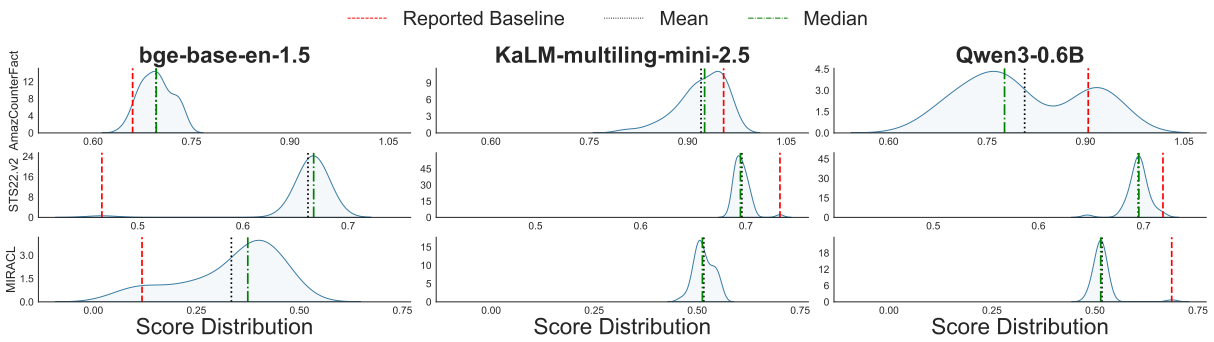


Figure 2: **Example of prompt deflation and inflation** on a sample of evaluated tasks and models. For all the scores and distributions see Appendix A. For KaLM model on MIRACL task, mteb does not have reported score.

prompt risks misrepresenting actual model performance. Furthermore, it is possible to overwrite the prompt within the models implementation – to allow customization to a specific prompt scheme – making it possible for the model developers to selectively choose the best performing prompts for a given task.

Current research on prompt sensitivity is focused mostly on generative LLMs (Wahle et al., 2024; Leidinger et al., 2023; Romanou et al., 2026; Mizrahi et al., 2024), while instruction-based embedding models remains largely underexplored. While recent efforts have addressed overfitting via paraphrases (Frank and Afli, 2026), no work systematically examines the effect of prompt variation on embedding evaluation. Our work bridges this gap by isolating the prompt as the primary variable to quantify its impact on reported performance.

In this paper, we question the validity of single-prompt evaluation for instruction-tuned embedding models. We investigate the following research questions:

(**RQ1**) To what extent are instruction-tuned embedding models sensitive to variations of prompts relevant to the task?
(**RQ2**) How do these variations influence the reported benchmark rankings?
(**RQ3**) To what degree is benchmark hacking facilitated by prompt hacking (Qiu et al., 2026)?

To answer these questions, we conduct an empirical study across 11 diverse tasks types including retrieval, classification, clustering, and semantic textual similarity (STS). For each task, we synthesize a set of 40 task-specific prompts with different LLMs, removed duplicates, and evaluate a suite of different open-weight embedding models. Our findings reveal that reported benchmark scores frequently misrepresent expected performance: the default prompt can act as an outlier, artificially deflating or inflating a model's perceived quality (Figure 1). More critically, we show that by adversarially selecting prompts, any model in our study can be made to appear as the top-ranked model on a simulated leaderboard.

We recommend that benchmarks transition from single-point estimates to multi-prompt evaluation — reporting score distributions or robustness metrics alongside point estimates — and that model developers either optimize for prompt robustness or adopt limited prompt schemes such as those used by the Jina v5 series (Akram et al., 2026).

2 Methodology

To test the influence of the prompts on the benchmark performance, we selected the following diverse list of tasks and datasets from `mteb`[2]: retrieval (MIRACLRetrievalHardNegatives.v2, Touche2020Retrieval.v3, FEVERHardNegatives (Zhang et al., 2023; Thakur et al., 2024; Thorne et al., 2018)), classification (TweetSentimentClassification, ImdbClassification, AmazonCounterfactualClassification (Barbieri et al., 2022; Maas et al., 2011; O'Neill et al., 2021)), clustering (MedrxivClusteringP2P.v2, StackExchangeClustering.v2 (Enevoldsen et al., 2025; Geigle et al., 2021)), and semantic similarity (STS15, STS14, STS22.v2 (Biçici, 2015; Bandhakavi et al., 2014; Chen et al., 2022)). We chose tasks either derived from MTEB (Muennighoff et al., 2022) or MMTEB (Enevoldsen et al., 2025). To allow for comparison to the English only BGE models we only evaluate on the English subsets.

[2]We use the dataset name in `mteb` to ensure reproducibility.

For each task, we generated a list of 15 prompts with gpt-oss-120b (OpenAI, 2025), Qwen3.6-35B-A3B (Qwen Team, 2026), and Gemma 4 26B (Team, 2026) with vLLM (Kwon et al., 2023) structured outputs with the default parameters. Models were instructed to generate prompts based on the combination of language, task description, and task name. A sample of prompts was manually evaluated to ensure that they were coherent and relevant for task (see Appendix D, and Appendix E for default prompts). The prompts generation was set up to resemble a real use-case, where the practitioner will rely on a custom, task-oriented prompt. We use synthetic prompt generation to reflect real practitioner behaviour and to avoid author bias in the reported effects. The final list of prompts was deduplicated based on Levenshtein distance and threshlod of 0.1. Number of prompts per task is provided in Table 1. The resulting prompts are short (mean 7.78 words, std 3.2), consistent with documented model usage examples. We validated that synthetic and human-written prompts produce largely comparable score distributions, and that where differences exist they do not affect our conclusions (see Appendix D).

Task	N
AmazonCounterfactualClassification	40
FEVERHardNegatives	39
STS22.v2	43
Touche2020Retrieval.v3	38
STS15	44
STS14	45
MIRACLRetrievalHardNegatives.v2	39
TweetSentimentClassification	42
MedrxivClusteringP2P.v2	39
StackExchangeClustering.v2	40
ImdbClassification	44

Table 1: Number of generated prompts used for each task after deduplication.

We seek to evaluate a broad selection including both multilingual models: Qwen3-Embedding-0.6B (Zhang et al., 2025), multilingual-e5-large-instruct (Wang et al., 2024), KaLM-embedding-multilingual-mini-instruct-v2.5 (Hu et al., 2025), as well as English models across a variety of sizes; bge-small-en-v1.5, bge-base-en-v1.5, and bge-large-en-v1.5 (Xiao et al., 2023). We ran the models on each task for each of the task-specific prompts, and compared the results with the reported scores on `mteb`. Models revisions and versions are available in Appendix B.

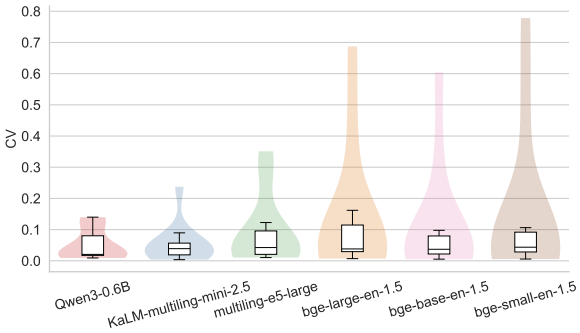


Figure 3: **Prompt sensitivity**: Coefficient of variance distributions across the models and prompts. We see that certain models have a high degree of variance across prompts.

3 Analysis and Discussion

3.1 Prompt Inflation and Deflation

We examine whether the reported single-prompt results are representative of a model's expected performance. In the Appendix A (Figure 7), the score distributions across the prompts for each model/task combination are presented, with the reported `mteb` score. In Appendix C, the analysis of confidence intervals is presented.

We observe two systematic patterns. First, several models exhibit *prompt deflation*: the reported score is much lower that what would be expected. For example, BGE models show prompt deflation on TweetSentimentClassification, AmazonCounterfactualClassification, and STS22.v2 tasks (see Figure 2), where its reported scores are in the lower tail of the distributions. In a few cases we even see the reported MTEB score is drastically below the main distribution. This suggests that the default prompts used for these models on these tasks are suboptimal and does not correspond to the expected performance.

Second, we observe *prompt inflation*: the reported score falls in the upper tail of the distribution, overstating expected performance. For example, the reported score of Qwen3 on MIRACL (see Figure 2) task aligns with a secondary mode in the upper range of the distribution, which is above the primary peak. This indicates that the default prompt produces an atypically strong result.

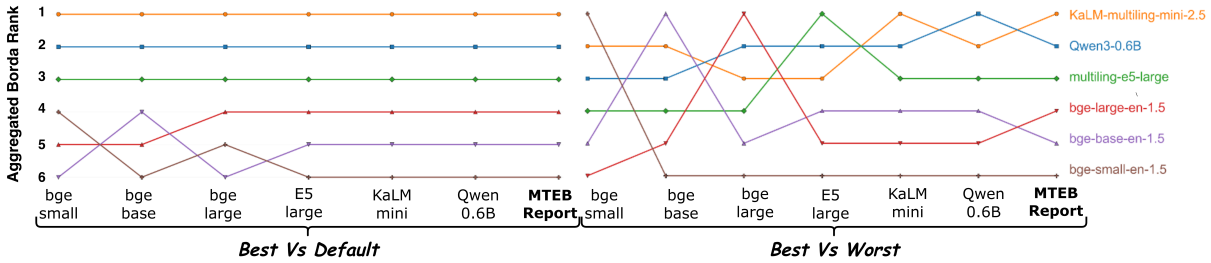


Figure 4: **Leaderboard sensitivity under adversarial prompt selection**. We show for each model two scenarios. Left: The model uses its best prompt while others use the default. Right: The model uses its best prompt while others use their worst. Ranks are computed via Borda scores across all tasks. In both cases we compare against the default configuration (MTEB Reported).

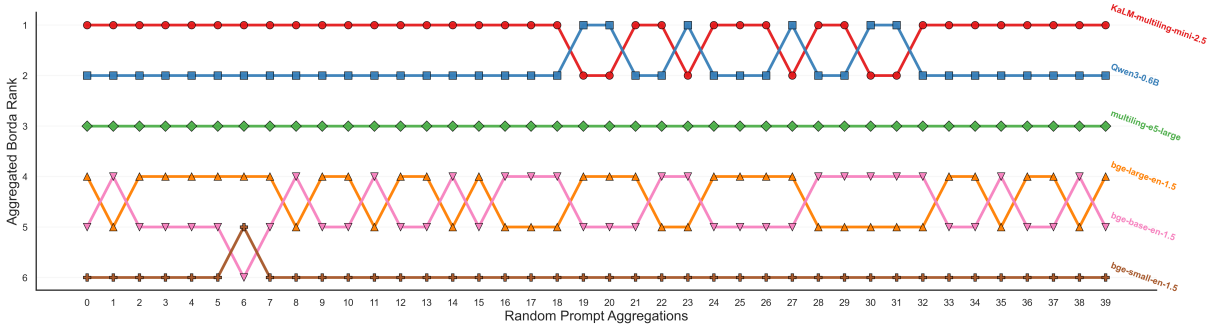


Figure 5: Leaderboard sensitivity under random prompt selection.

3.2 Prompt Sensitivity Analysis

To answer RQ1, we analyze distributions of scores and coefficients of variation (CVs) of the scores obtained by the tasks per model (see Figure 3). The higher the CV is, the less robust the model is for variations in the prompt. We chose CV as it is unit-free, which allows us to compare scores for different tasks. While median CVs are moderate (2–6%), the distributions exhibit long tails: for the BGE models, individual task CVs reach 30–45%, confirming that specific model–task combinations are highly sensitive to prompt choice.

3.3 Effect on Leaderboard Rankings

Following RQ2, we investigate whether prompt sensitivity can be exploited to alter model rankings. Since `mteb` allows model contributors to specify custom prompts, a model developer could select a prompt that maximizes their model's score. To simulate this, we construct adversarial leaderboard configurations: for each target model, we assign it its best-performing prompt across tasks while assigning all other models either their worst-performing or default prompts, and compute the aggregated Borda rank. The rank changes per configuration are presented on Figure 4.

Under default prompts, KaLM and Qwen3 occupy ranks 1–2, e5-large rank 3, and the BGE variants ranks 4–6. A well-meaning developer who simply reports the prompt that worked best during development can therefore distort the leaderboard. However, under adversarial selection (an upper bound on manipulability), every model in our study can be promoted to rank 1, including bge-small, which rises from last place. This configuration does not model public submissions, where each contributor sets only their own prompt; rather, it corresponds to a setting where one party controls all prompt assignments, such as a selective write-up or public release, and it bounds how far favorable prompt selection can move a ranking.

Importantly, the milder scenario with the best performing prompt for the target model and default for others, still produces meaningful rank changes. For example, bge-large rises from rank 5 to rank 3 under this configuration, overtaking e5-large. This demonstrates that even a single model contributor optimizing their prompt, without any adversarial intent toward competitors, can distort the leaderboard.

We run the leaderboard sensitivity analysis sampling random prompts as well. The results are present in Figure 5. They support our prior findings and display similar trends when it comes to

the rankings.

3.4 Prompt Inflation in Reported Scores

Lastly, to examine whether some models systematically overfit (RQ3), we analyzed how the reported MTEB scores correspond to an observed score distribution. For each model–task combination, we calculated the empirical probability of the reported score across obtained scores (see Figure 6). Our results reveal a systematic issue: most models report scores above the median of the prompt distribution, potentially of the intended instruction. It can be interpreted as a systematic gap between the single reported estimate and the performance expected across plausible prompts. This pattern is analogous to p-hacking (Simmons et al., 2011) in that it exploits undisclosed degrees of freedom in evaluation without constituting an explicit violation of submission policies — we term this *prompt hacking* (Qiu et al., 2026): the selective reporting of a single favorable prompt while remaining compliant, without modifying model weights, training data, or architecture. Importantly, prompt hacking need not be intentional: just as analytical bias can arise from well-meaning researchers exploring analytical choices without recognizing the inflationary effect (Gelman and Loken, 2016), a model developer may simply report the prompt that worked best during development without deliberate intent to game the benchmark.

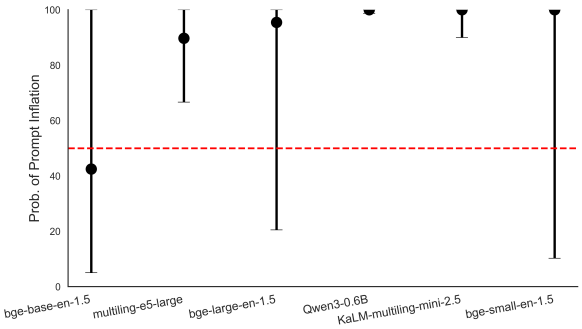


Figure 6: **Probability of prompt inflation**, i.e. reported score higher than the median, across tasks. The uncertainty is a 95% bootstrapped confidence interval.

4 Conclusion

In this paper, we conducted an empirical study of prompt sensitivity across 6 models and 11 tasks. We find that reported benchmark scores frequently misrepresent expected performance: the MTEB default prompt can systematically overstate or understate a model's quality. Under adversarial prompt selection, every model in our study can be promoted to rank 1, including models that rank last under default evaluation. Even a single developer optimizing their prompt can shift rankings by multiple positions. We recommend that benchmarks adopt multi-prompt evaluation or similar strategies to ensure that the reported rankings faithfully reflect the models' quality.

Limitations

Model scope: We evaluate 6 open-weight models with up to 0.6B parameters. Larger instruction-tuned embedding models and closed-source models (e.g., OpenAI text-embedding-3, Cohere embed-v4) may exhibit different sensitivity profiles.

Prompt generation and sample size: Our prompt sets consist of ∼40 synthetically generated prompts per task, produced by language models. This is a relatively small sample: variance estimates (e.g., CV) computed over ∼40 points carry non-trivial uncertainty, and the synthetic generation process may not fully capture the diversity of prompts that practitioners write in practice. A larger, more diverse prompt set would yield more robust conclusions.

Language coverage: We restrict our analysis to English. Multilingual settings may introduce additional sources of prompt sensitivity related to language-specific phrasing conventions and translation artifacts.

Prompt variability as a practitioner asset: Our study frames prompt sensitivity primarily as a threat to benchmark validity, but the same variability can be an asset in practice. A developer who tries even a small set of task-specific prompts — as our results show — can gain up to 10 percentage points in performance at essentially no cost. This work does not evaluate strategies for exploiting this, such as prompt search or ensembling, and the optimal number of prompts to try in a practical setting remains an open question.

Acknowledgments

Yevhen Kostiuk and Kenneth Enevoldsen are funded by the Danish Foundation Models project (4378-00001B). Kenneth Enevoldsen is additionally funded by the European Union, Horizon Europe (101178170), the Danish National Research Foundation (DNRF193), the Aage and Johanne Louis-Hansens Foundation (25-1-17733), and the Augustinus Foundation (2025-0299).

Part of the computation done for this project was performed on the UCloud interactive HPC system, which is managed by the eScience Center at the University of Southern Denmark.

References

Mohammad Kalim Akram, Saba Sturua, Nastia Havriushenko, Quentin Herreros, Michael Günther, Maximilian Werk, and Han Xiao. 2026. jina-embeddings-v5-text: Task-targeted embedding distillation. *arXiv preprint arXiv:2602.15547*.

Anthropic. 2026. Claude. `https://claude.ai/`. [Large language model].

Anil Bandhakavi, Nirmalie Wiratunga, Deepak P, and Stewart Massie. 2014. Generating a word-emotion lexicon from #emotional tweets. In *Proceedings of the Third Joint Conference on Lexical and Computational Semantics (*SEM 2014)*, pages 12–21, Dublin, Ireland. Association for Computational Linguistics and Dublin City University.

Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 258–266, Marseille, France. European Language Resources Association.

Ergun Biçici. 2015. RTM-DCU: Predicting semantic similarity with referential translation machines. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 56–63, Denver, Colorado. Association for Computational Linguistics.

Xi Chen, Ali Zeynali, Chico Camargo, Fabian Flöck, Devin Gaffney, Przemyslaw Grabowicz, Scott Hale, David Jurgens, and Mattia Samory. 2022. SemEval-2022 task 8: Multilingual news article similarity. In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 1094–1106, Seattle, United States. Association for Computational Linguistics.

Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Siblini, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Rystrøm, Roman Solomatin, and 67 others. 2025. Mmteb: Massive multilingual text embedding benchmark. *arXiv preprint arXiv:2502.13595*.

Manuel Frank and Haithem Afli. 2026. PTEB: Towards robust text embedding evaluation via stochastic paraphrasing at evaluation time with LLMs. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2832–2851, Rabat, Morocco. Association for Computational Linguistics.

Gregor Geigle, Nils Reimers, Andreas Rücklé, and Iryna Gurevych. 2021. Tweac: Transformer with extendable qa agent classifiers. *arXiv preprint*, abs/2104.07081.

Andrew Gelman and Eric Loken. 2016. The statistical crisis in science. *The best writing on mathematics (Pitici M, ed)*, 102:305–318.

Xinshuo Hu, Zifei Shan, Xinping Zhao, Zetian Sun, Zhenyu Liu, Dongfang Li, Shaolin Ye, Xinyuan Wei, Qian Chen, Baotian Hu, Haofen Wang, Jun Yu, and Min Zhang. 2025. Kalm-embedding: Superior training data brings a stronger embedding model. *Preprint*, arXiv:2501.01028.

Ehsan Kamalloo, Nandan Thakur, Carlos Lassance, Xueguang Ma, Jheng-Hong Yang, and Jimmy Lin. 2024. Resources for brewing beir: Reproducible reference models and statistical analyses. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '24, page 1431–1440, New York, NY, USA. Association for Computing Machinery.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Alina Leidinger, Robert van Rooij, and Ekaterina Shutova. 2023. The language of prompting: What linguistic properties make a prompt successful? *Preprint*, arXiv:2311.01967.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland,

Oregon, USA. Association for Computational Linguistics.

Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2024. State of what art? a call for multi-prompt LLM evaluation. *Transactions of the Association for Computational Linguistics*, 12:933–949.

Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. Mteb: Massive text embedding benchmark. *arXiv preprint arXiv:2210.07316*.

James O'Neill, Polina Rozenshtein, Ryuichi Kiryo, Motoko Kubota, and Danushka Bollegala. 2021. I wish I would have loved this one, but I didn't – a multilingual dataset for counterfactual detection in product review. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7092–7108, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

OpenAI. 2025. gpt-oss-120b & gpt-oss-20b model card. *Preprint*, arXiv:2508.10925.

Xiaoyun Qiu, Yang Yu, and Haifeng Xu. 2026. On benchmark hacking in ml contests: Modeling, insights and design. *arXiv preprint arXiv:2604.22230*.

Qwen Team. 2026. Qwen3.6-35B-A3B: Agentic coding power, now open to all.

Angelika Romanou, Mark Ibrahim, Candace Ross, Chantal Shaib, Kerem Oktar, Samuel J. Bell, Anaelia Ovalle, Jesse Dodge, Antoine Bosselut, Koustuv Sinha, and Adina Williams. 2026. Brittlebench: Quantifying llm robustness via prompt sensitivity. *Preprint*, arXiv:2603.13285.

Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. 2011. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11):1359–1366.

Aaditya Singh, Adam Fry, and et al. 2025. Openai gpt-5 system card. *Preprint*, arXiv:2601.03267.

Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2023. One embedder, any task: Instruction-finetuned text embeddings. *Preprint*, arXiv:2212.09741.

Gemma Team. 2026. Gemma 4 technical report. *Preprint*, arXiv:2607.02770.

Nandan Thakur, Luiz Bonifacio, Maik Fröbe, Alexander Bondarenko, Ehsan Kamalloo, Martin Potthast, Matthias Hagen, and Jimmy Lin. 2024. Systematic evaluation of neural retrieval models on the Touché 2020 argument retrieval subset of BEIR. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*.

Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.

Jan Philip Wahle, Terry Ruas, Yang Xu, and Bela Gipp. 2024. Paraphrase types elicit prompt engineering capabilities. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11004–11033, Miami, Florida, USA. Association for Computational Linguistics.

Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024. Multilingual e5 text embeddings: A technical report. *arXiv preprint arXiv:2402.05672*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. Huggingface's transformers: State-of-the-art natural language processing. *Preprint*, arXiv:1910.03771.

Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. C-pack: Packaged resources to advance general chinese embedding. *Preprint*, arXiv:2309.07597.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2026. A-mem: Agentic memory for llm agents. *Advances in Neural Information Processing Systems*, 38:17577–17604.

Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023. MIRACL: A Multilingual Retrieval Dataset Covering 18 Diverse Languages. *Transactions of the Association for Computational Linguistics*, 11:1114–1131.

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.

A Detailed Scores

B Model Revisions

For our experiments, we used the following model and revisions from HuggingFace Hub (see Table 2) (Wolf et al., 2020) and MTEB version 2.12.30.

C Confidence Intervals for Borda Ranks

To assess the sensitivity of the MTEB leaderboard to prompt selection, we recompute the Borda-count aggregate ranking and quantify the resulting rank uncertainty via a nonparametric bootstrap. To estimate confidence intervals (CIs) on the resulting ranks, we bootstrap over two variations: (1) the set of evaluation tasks, resampled with replacement to reflect uncertainty in benchmark composition, and (2) the set of sampled prompts per task-model pair, resampled with replacement (excluding the fixed default prompt) to reflect the fact that the observed best/worst scores are themselves noisy extreme-value estimates drawn from a finite prompt sample. We repeat this joint resampling procedure 1,000 times, recomputing the full Borda ranking on each replicate, and report the mean rank together with the 5th and 95th percentiles of the bootstrap rank distribution as a 90% confidence interval for each model under each regime.

The results are presented in Figure 8, **??**, and Table 4. Under the reported (reportedprompt) leaderboard, models separate into three loose tiers — KaLM-multiling-mini-2.5 and Qwen3-0.6B occupy the top two positions (mean ranks 1.46 and 1.54, 90% CI [1,2] for both), multiling-e5-large sits alone in third (mean rank 3.01, CI [3,3]), and bge-large-en-1.5, bge-base-en-1.5, and bge-small-en-1.5 trail behind, with bge-base and bge-small's intervals overlapping ([5,6] and [4,6] respectively). The overlapping intervals at both the top and bottom of the table indicate that the reported ordering within these pairs is not statistically well-separated. This finding futher supports the claim that ordering is highly influenced by prompt selection: when a model is granted its best prompt while competitors are forced to use their worst, every one of the six models is able to reach or approach rank 1. Notably, KaLM-multiling-mini-2.5 and Qwen3-0.6B are the only two models that reach rank 1 with zero-width confidence intervals ([1,1]) under their respective best-self/worst-others regimes, suggesting their advantage over the remaining four models is comparatively robust to prompt choice even though the top-two ordering between themselves is not. In contrast, the best prompt for one model against competitors' reported scores produces far smaller rank shifts for larger models (KaLM and Qwen3's), showing that most of the observed leaderboard instability stems from the choice of prompt used for competing models rather than from the model's own prompt optimization alone.

D Prompt Evaluation

We manually evaluated 100 generated prompts. For each prompt, we looked at the task and task description, and assigned one of the following labels:

- Acceptable. The prompt sounds natural, relevant to the task, and the task is clearly stated.
- Irrelevant. The prompt is irrelevant to the task.
- Unnatural. The prompt does not sound correct and is not natural.

Based on our analysis *all* the generated prompts were considered to be acceptable. We noticed that some prompts contained name of the task in them, e.g. "Similarity score request, eng, SemEval STS 2014:". We marked these prompts as acceptable, as such prompts contain reasonable instructions, and the task name may provide an indication of the task to the model. In the case of the example, SemEval is a well known shared task, therefore the model may potentially know its relation towards STS.

To show that our findings are generalizable to natural human prompts, we created 10 prompts per task for BGE family of models and compared the human prompts performance with the machine generated prompts. The results are presented on Figure 9 and Figure 10. The distributions were tested with the permutation test on similarity of distributions.

The figure shows that the human and synthetic prompts perform similarly in most cases. In the cases, where the test did not reject the hypothesis of distributional similarity, the analysis of the reported score and performance distribution does not contradict findings of the paper.

The acceptability check was carried out by a single annotator (one of the authors). A limitation of this approach is that a single annotator may hold an implicit bias toward accepting the generated prompts. The scheme used is a coarse coherence filter (Acceptable / Irrelevant / Unnatural) that targets clearly broken instructions rather than fine-grained

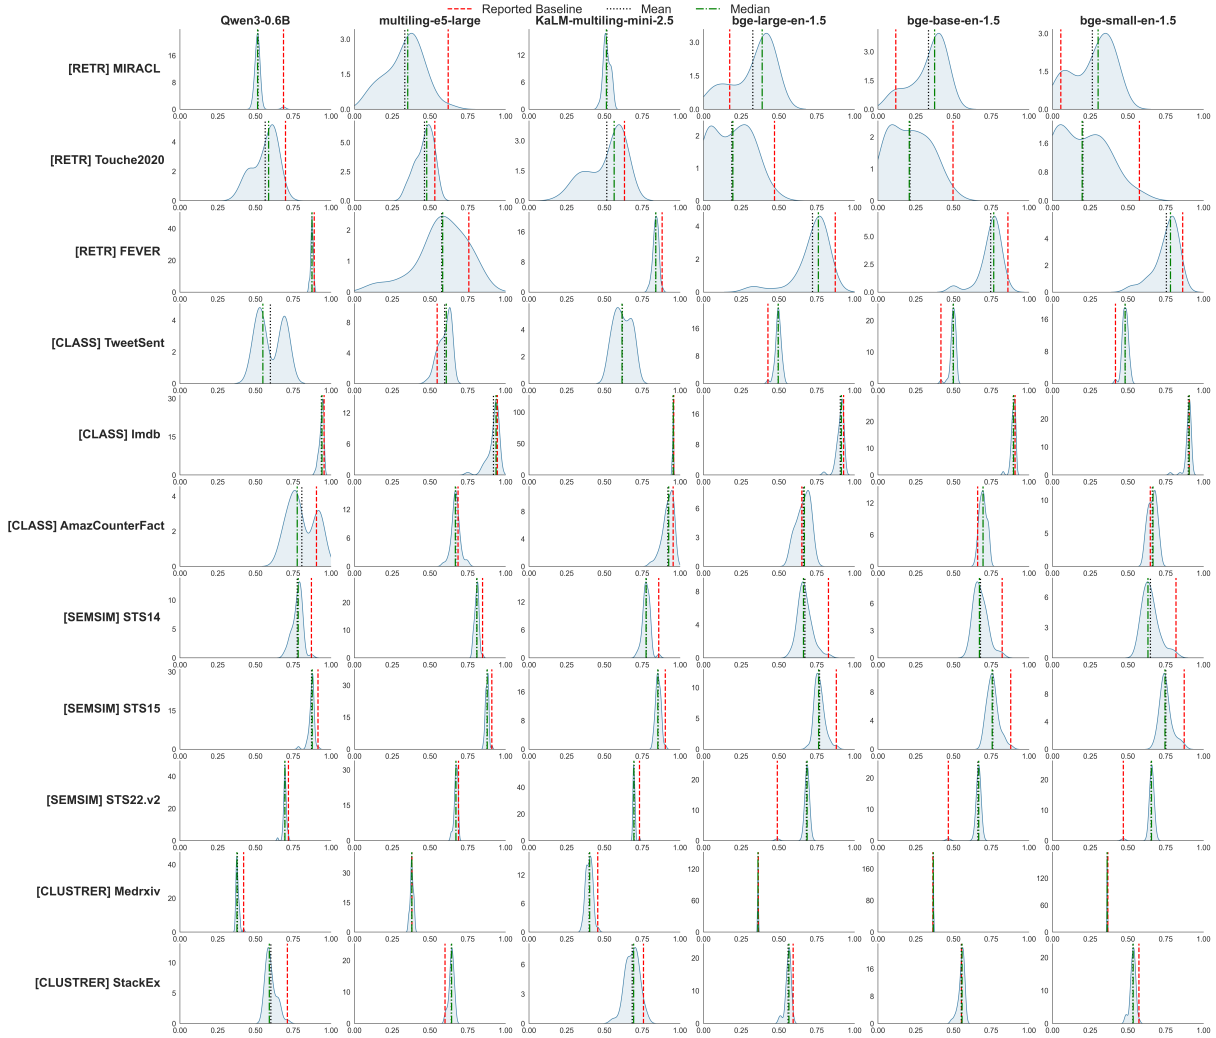


Figure 7: Detailed scores per model, metric, and task. For some tasks MTEB does not currently report a score hence the red line is missing.

Model	Commit
Qwen/Qwen3-Embedding-0.6B	97b0c614
intfloat/multilingual-e5-large-instruct	274baa43
KaLM-Embedding/KaLM-embedding-multilingual-mini-instruct-v2.5	52c687bb
BAAI/bge-small-en-v1.5	5c38ec7c
BAAI/bge-base-en-v1.5	a5beb1e3
BAAI/bge-large-en-v1.5	d4aa6901

Table 2: Model revisions used in experiments.

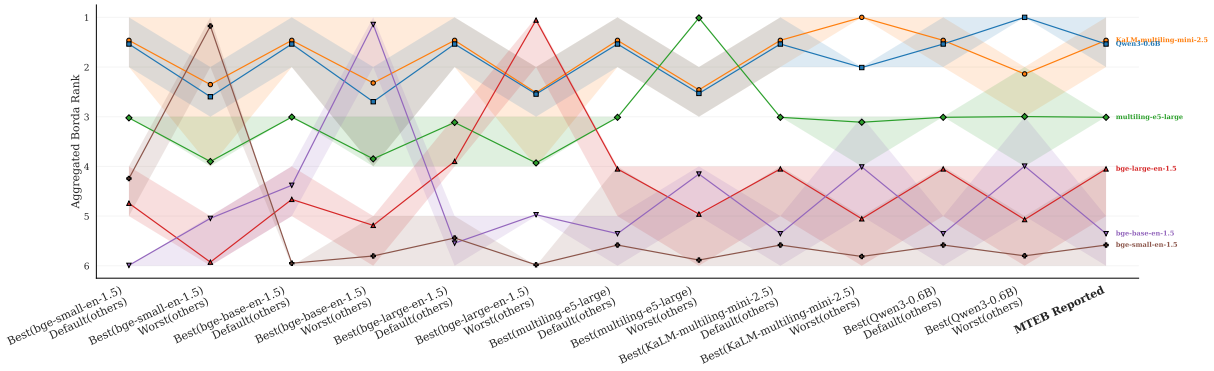


Figure 8: CIs of the ranks.

Model	Cond.	Rank	Borda
Qwen3-0.6B	S-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	S-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	S-D	3.023 [3–3]	2.636 [2.09–3.09]
bge-large-en-1.5	S-D	4.741 [4–5]	1.281 [0.82–1.73]
bge-base-en-1.5	S-D	5.992 [6–6]	0.540 [0.27–0.82]
bge-small-en-1.5	S-D	4.244 [4–5]	1.633 [1.18–2.09]
Qwen3-0.6B	S-W	2.598 [2–3]	3.283 [2.82–3.73]
KaLM	S-W	2.353 [1–4]	3.371 [2.56–4.20]
multiling-e5-large	S-W	3.902 [3–4]	2.273 [1.55–3.00]
bge-large-en-1.5	S-W	5.927 [5–6]	0.554 [0.27–0.82]
bge-base-en-1.5	S-W	5.046 [5–6]	1.105 [0.64–1.55]
bge-small-en-1.5	S-W	1.174 [1–2]	4.065 [3.64–4.45]
Qwen3-0.6B	B-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	B-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	B-D	3.007 [3–3]	2.636 [2.09–3.09]
bge-large-en-1.5	B-D	4.664 [4–5]	1.371 [1.00–1.82]
bge-base-en-1.5	B-D	4.380 [4–5]	1.538 [1.27–1.82]
bge-small-en-1.5	B-D	5.949 [6–6]	0.545 [0.09–1.09]
Qwen3-0.6B	B-W	2.698 [2–4]	3.192 [2.78–3.60]
KaLM	B-W	2.321 [1–4]	3.371 [2.56–4.20]
multiling-e5-large	B-W	3.848 [3–4]	2.346 [1.64–3.00]
bge-large-en-1.5	B-W	5.187 [5–6]	0.919 [0.73–1.09]
bge-base-en-1.5	B-W	1.143 [1–2]	4.149 [3.73–4.55]
bge-small-en-1.5	B-W	5.803 [5–6]	0.666 [0.18–1.18]

Table 3: Model rankings and Borda scores for BGE-small and BGE-base prompt conditions.

quality, which limits the room for subjective disagreement. The annotation results excluded nothing: all prompts were judged Acceptable, so no regeneration or replacement was performed and no borderline decision affected the prompt set. Had any prompt been rejected, our policy was to remove rejected prompts for that task to regenerate the same number of missing prompts and evaluate them again, until a full set of acceptable prompts was obtained.

E Default MTEB Prompts

In Table 5, Table 6, and Table 7 the default MTEB prompts are present per each model–task combination.

F Best and Worst Performing Prompts

G AI Tools Usage

In this work, we utilized AI tools to assist with code generation, debugging, and spell-checking/grammatical editing of the manuscript text. Specifically, we used Anthropic's Claude models (Anthropic, 2026) and OpenAI's ChatGPT (Singh et al., 2025).

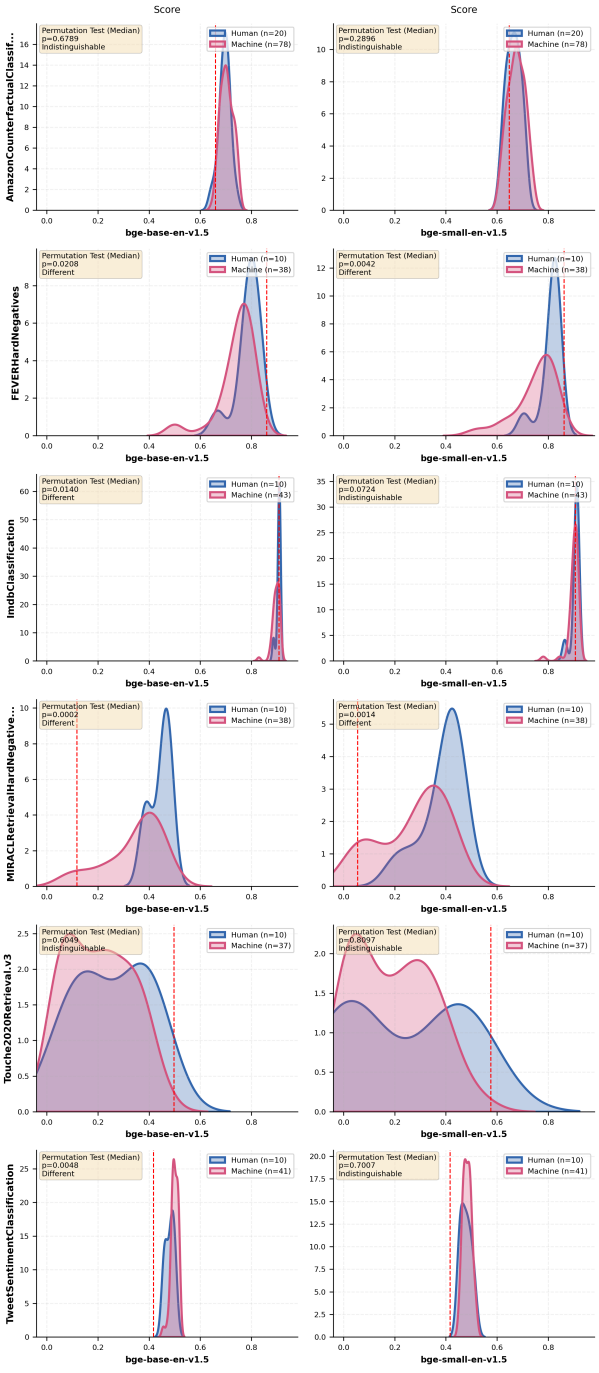


Figure 9: Comparison of human vs machine generated prompt score distribution for Classification and Retrieval tasks. Red dashed line indicates a reported score on MTEB.

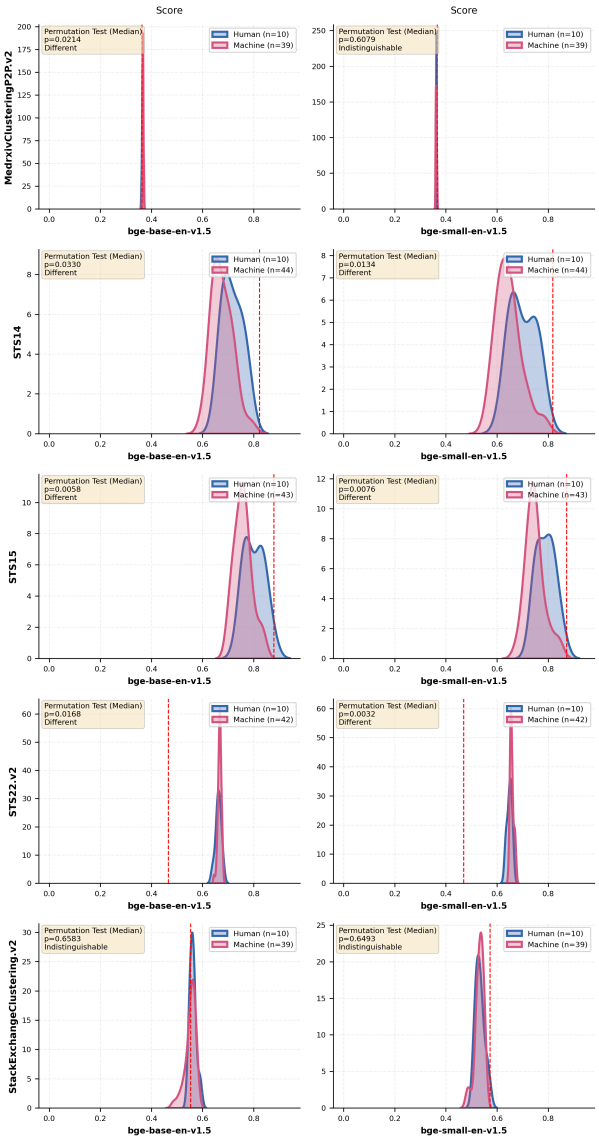


Figure 10: Comparison of human vs machine generated prompt score distribution for Clustering and STS tasks. Red dashed line indicates a reported score on MTEB.

Model	Cond.	Rank	Borda
Qwen3-0.6B	L-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	L-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	L-D	3.115 [3–4]	2.540 [2.00–3.00]
bge-large-en-1.5	L-D	3.899 [3–4]	2.006 [1.55–2.45]
bge-base-en-1.5	L-D	5.545 [5–6]	0.724 [0.45–0.91]
bge-small-en-1.5	L-D	5.441 [5–6]	0.820 [0.36–1.36]
Qwen3-0.6B	L-W	2.545 [2–3]	3.237 [2.80–3.70]
KaLM	L-W	2.517 [2–4]	3.232 [2.44–4.00]
multiling-e5-large	L-W	3.930 [3–4]	2.145 [1.45–2.91]
bge-large-en-1.5	L-W	1.055 [1–2]	4.232 [3.82–4.64]
bge-base-en-1.5	L-W	4.972 [5–5]	1.197 [0.82–1.55]
bge-small-en-1.5	L-W	5.981 [6–6]	0.581 [0.27–0.91]
Qwen3-0.6B	E-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	E-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	E-D	3.012 [3–3]	2.726 [2.18–3.18]
bge-large-en-1.5	E-D	4.051 [4–5]	1.643 [1.18–2.09]
bge-base-en-1.5	E-D	5.354 [5–6]	0.901 [0.64–1.09]
bge-small-en-1.5	E-D	5.583 [4–6]	0.820 [0.36–1.36]
Qwen3-0.6B	E-W	2.531 [2–3]	3.450 [3.11–3.80]
KaLM	E-W	2.459 [2–3]	3.481 [2.78–4.00]
multiling-e5-large	E-W	1.011 [1–1]	4.544 [4.18–4.82]
bge-large-en-1.5	E-W	4.960 [4–6]	1.064 [0.73–1.45]
bge-base-en-1.5	E-W	4.154 [4–5]	1.527 [1.09–1.91]
bge-small-en-1.5	E-W	5.885 [5–6]	0.619 [0.27–1.00]
Qwen3-0.6B	K-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	K-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	K-D	3.012 [3–3]	2.726 [2.18–3.18]
bge-large-en-1.5	K-D	4.051 [4–5]	1.643 [1.18–2.09]
bge-base-en-1.5	K-D	5.354 [5–6]	0.901 [0.64–1.09]
bge-small-en-1.5	K-D	5.583 [4–6]	0.820 [0.36–1.36]
Qwen3-0.6B	K-W	2.010 [2–2]	3.693 [3.44–3.91]
KaLM	K-W	1.000 [1–1]	5.000 [5.00–5.00]
multiling-e5-large	K-W	3.111 [3–4]	2.479 [1.82–3.09]
bge-large-en-1.5	K-W	5.055 [4–6]	1.182 [0.82–1.55]
bge-base-en-1.5	K-W	4.011 [3–5]	1.800 [1.36–2.19]
bge-small-en-1.5	K-W	5.813 [5–6]	0.818 [0.36–1.36]
Qwen3-0.6B	Q-D	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	Q-D	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	Q-D	3.012 [3–3]	2.726 [2.18–3.18]
bge-large-en-1.5	Q-D	4.051 [4–5]	1.643 [1.18–2.09]
bge-base-en-1.5	Q-D	5.354 [5–6]	0.901 [0.64–1.09]
bge-small-en-1.5	Q-D	5.583 [4–6]	0.820 [0.36–1.36]
Qwen3-0.6B	Q-W	1.000 [1–1]	4.902 [4.73–5.00]
KaLM	Q-W	2.139 [2–3]	3.269 [2.64–3.75]
multiling-e5-large	Q-W	2.997 [2–4]	2.592 [1.91–3.18]
bge-large-en-1.5	Q-W	5.070 [4–6]	1.240 [0.91–1.55]
bge-base-en-1.5	Q-W	3.994 [3–5]	1.876 [1.36–2.36]
bge-small-en-1.5	Q-W	5.800 [5–6]	0.895 [0.36–1.45]
Qwen3-0.6B	MTEB	1.537 [1–2]	4.397 [4.13–4.64]
KaLM	MTEB	1.463 [1–2]	4.451 [4.10–4.80]
multiling-e5-large	MTEB	3.012 [3–3]	2.726 [2.18–3.18]
bge-large-en-1.5	MTEB	4.051 [4–5]	1.643 [1.18–2.09]
bge-base-en-1.5	MTEB	5.354 [5–6]	0.901 [0.64–1.09]
bge-small-en-1.5	MTEB	5.583 [4–6]	0.820 [0.36–1.36]

Table 4: Model rankings and Borda scores for the remaining prompt conditions.

Table 5: Default prompts by model and task.

Model	Task	Prompt
Qwen/Qwen3-Embedding-0.6B	AmazonCounterfactual	Classify a given Amazon customer review text as either counterfactual or not-counterfactual
Qwen/Qwen3-Embedding-0.6B	FEVERHardNegatives	Retrieve text based on user query.
Qwen/Qwen3-Embedding-0.6B	STS22.v2	Retrieve semantically similar text.
Qwen/Qwen3-Embedding-0.6B	Touche2020Retrieval.v3	Retrieve text based on user query.
Qwen/Qwen3-Embedding-0.6B	STS15	Retrieve semantically similar text.
Qwen/Qwen3-Embedding-0.6B	STS14	Retrieve semantically similar text.
Qwen/Qwen3-Embedding-0.6B	MIRACLRetrievalHard	Given a question, retrieve Wikipedia passages that answer the question
Qwen/Qwen3-Embedding-0.6B	TweetSentiment	Classify user passages.
Qwen/Qwen3-Embedding-0.6B	MedrxivClusteringP2P.v2	Identify the main category of Medrxiv papers based on the titles and abstracts
Qwen/Qwen3-Embedding-0.6B	StackExchangeClustering.v	Identify the topic or theme of StackExchange posts based on the titles
Qwen/Qwen3-Embedding-0.6B	ImdbClassification	Classify the sentiment expressed in the given movie review text from the IMDB dataset
BAAI/bge-small-en-v1.5	AmazonCounterfactual	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	FEVERHardNegatives	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	STS22.v2	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	Touche2020Retrieval.v3	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	STS15	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	STS14	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	MIRACLRetrievalHard	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	TweetSentiment	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	MedrxivClusteringP2P.v2	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	StackExchangeClustering	Represent this sentence for searching relevant passages:
BAAI/bge-small-en-v1.5	ImdbClassification	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	AmazonCounterfactual	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	FEVERHardNegatives	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	STS22.v2	Represent this sentence for searching relevant passages:

Table 6: Default prompts by model and task.

Model	Task	Prompt
BAAI/bge-base-en-v1.5	Touche2020Retrieval.v3	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	STS15	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	STS14	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	MIRACLRetrievalHard	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	TweetSentiment	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	MedrxivClusteringP2P.v2	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	StackExchangeClustering.v	Represent this sentence for searching relevant passages:
BAAI/bge-base-en-v1.5	ImdbClassification	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	AmazonCounterfactual	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	FEVERHardNegatives	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	STS22.v2	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	Touche2020Retrieval.v3	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	STS15	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	STS14	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	MIRACLRetrievalHard	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	TweetSentiment	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	MedrxivClusteringP2P.v2	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	StackExchangeClustering.v	Represent this sentence for searching relevant passages:
BAAI/bge-large-en-v1.5	ImdbClassification	Represent this sentence for searching relevant passages:
intfloat/multilingual-e5-	AmazonCounterf	Classify a given Amazon customer review text as either counterfactual or not-counterfactual
intfloat/multilingual-e5-	FEVERHardNegatives	Retrieve text based on user query.
intfloat/multilingual-e5-	STS22.v2	Retrieve semantically similar text.
intfloat/multilingual-e5-	Touche2020Retrieval.v3	Retrieve text based on user query.
intfloat/multilingual-e5-	STS15	Retrieve semantically similar text.
intfloat/multilingual-e5-	STS14	Retrieve semantically similar text.
intfloat/multilingual-e5-	MIRACLRetrievalHard	Given a question, retrieve Wikipedia passages that answer the question
intfloat/multilingual-e5-	TweetSentiment	Classify user passages.
intfloat/multilingual-e5-	MedrxivClusteringP2P.v2	Identify the main category of Medrxiv papers based on the titles and abstracts
intfloat/multilingual-e5-	StackExchangeClustering	Identify the topic or theme of StackExchange posts based on the titles
intfloat/multilingual-e5-	ImdbClassification	Classify the sentiment expressed in the given movie review text from the IMDB dataset
KaLM-Embedding/KaLM-embed	AmazonCounterfactual	Classify a given Amazon customer review text as either counterfactual or not-counterfactual

Table 7: Default prompts by model and task.

Model	Task	Prompt
KaLM-Embedding/KaLM-embed	FEVERHardNegatives	Retrieve text based on user query.
KaLM-Embedding/KaLM-embed	STS22.v2	Retrieve semantically similar text.
KaLM-Embedding/KaLM-embed	Touche2020Retrieval.v3	Retrieve text based on user query.
KaLM-Embedding/KaLM-embed	STS15	Retrieve semantically similar text
KaLM-Embedding/KaLM-embed	STS14	Retrieve semantically similar text
KaLM-Embedding/KaLM-embed	MIRACLRetrievalHard	Given a question, retrieve Wikipedia passages that answer the question
KaLM-Embedding/KaLM-embed	TweetSentimentClassificat	Classify user passages.
KaLM-Embedding/KaLM-embed	MedrxivClusteringP2P.v2	Identify the main category of Medrxiv papers based on the titles and abstracts
KaLM-Embedding/KaLM-embed	StackExchangeClustering.v	Identify the topic or theme of StackExchange posts based on the titles
KaLM-Embedding/KaLM-embed	ImdbClassification	Classify the sentiment expressed in the given movie review text from the IMDB dataset

Table 8: Best and worst prompts for task: AmazCounterFact

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.958	Category prediction for Amazon review counterfactual detection pair in eng:
Qwen3-0.6B	Worst	0.808	Assign a category to this text:
	Best	0.931	Classify this Amazon review pair for counterfactual detection in eng:
bge-base-en-1.5	Worst	0.67	Assign a classification label to these Amazon reviews in eng:
	Best	0.734	Label this example:
bge-large-en-1.5	Worst	0.648	Category prediction for Amazon review counterfactual detection pair in eng:
	Best	0.737	Classify:
bge-small-en-1.5	Worst	0.571	Counterfactual detection pair classification for Amazon customer reviews (eng):
	Best	0.711	Classify:
multiling-e5-large	Worst	0.61	Counterfactual detection pair classification for Amazon customer reviews (eng):
	Best	0.75	Assign a counterfactual detection label:
	Worst	0.589	Predict the class:

Table 9: Best and worst prompts for task: FEVERHardNegatives

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.882	default
Qwen3-0.6B	Worst	0.79	Relevant passage retrieval in eng (FEVER dataset):
	Best	0.89	Retrieval task prompt:
bge-base-en-1.5	Worst	0.853	Encode this claim for retrieval:
	Best	0.86	default
bge-large-en-1.5	Worst	0.5	Text representation for retrieval:
	Best	0.873	default
bge-small-en-1.5	Worst	0.309	Text representation for retrieval:
	Best	0.862	default
multiling-e5-large	Worst	0.507	Text representation for retrieval:
	Best	0.832	Retrieve relevant documents in eng for the FEVER claim:
	Worst	0.131	Encode this text for retrieval:

Table 10: Best and worst prompts for task: ImdbClassification

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.959	default
Qwen3-0.6B	Worst	0.945	Determine the class for the given passage:
	Best	0.954	default
bge-base-en-1.5	Worst	0.894	Determine the class for the given passage:
	Best	0.914	Assign a label to this review:
bge-large-en-1.5	Worst	0.829	Passage: eng - Classify this example from the Large Movie Review Dataset:
	Best	0.935	Assign a label to this review:
bge-small-en-1.5	Worst	0.798	Passage: eng - Classify this example from the Large Movie Review Dataset:
	Best	0.917	Assign a label to this review:
multiling-e5-large	Worst	0.78	Passage: eng - Classify this example from the Large Movie Review Dataset:
	Best	0.95	Classify the movie review:
	Worst	0.75	Determine the class for the given passage:

Table 11: Best and worst prompts for task: MIRACLRetrievalHardNegatives.v2

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.556	Search for documents using this query:
Qwen3-0.6B	Worst	0.461	Instruction for MIRACL retrieval task:
	Best	0.686	default
bge-base-en-1.5	Worst	0.485	Text to encode for multilingual retrieval in eng:
	Best	0.481	Search query:
bge-large-en-1.5	Worst	0.076	Text to encode for multilingual retrieval in eng:
	Best	0.475	Search query:
bge-small-en-1.5	Worst	0.048	Query representation for retrieval:
	Best	0.455	Search query:
multiling-e5-large	Worst	0.019	Encode query for multilingual retrieval:
	Best	0.621	default
	Worst	0.089	Encode this text for retrieval:

Table 12: Best and worst prompts for task: MedrxivClusteringP2P.v2

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.456	default
Qwen3-0.6B	Worst	0.355	Group similar texts:
	Best	0.422	default
bge-base-en-1.5	Worst	0.368	Medrxiv title+abstract clustering task:
	Best	0.371	Cluster representation:
bge-large-en-1.5	Worst	0.362	Text clustering input:
	Best	0.367	Clustering input:
bge-small-en-1.5	Worst	0.356	medrxiv titles+abstract clustering input:
	Best	0.367	default
multiling-e5-large	Worst	0.359	Text clustering input:
	Best	0.399	Clustering of titles+abstract from medrxiv across 51 categories:
	Worst	0.356	Group this text:

Table 13: Best and worst prompts for task: STS14

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.86	default
Qwen3-0.6B	Worst	0.713	Measure semantic similarity for eng inputs from SemEval STS 2014 dataset:
	Best	0.871	default
bge-base-en-1.5	Worst	0.692	Text pair similarity assessment (SemEval STS 2014, eng):
	Best	0.823	default
bge-large-en-1.5	Worst	0.597	Text pair similarity assessment (SemEval STS 2014, eng):
	Best	0.828	default
bge-small-en-1.5	Worst	0.582	Text pair similarity assessment:
	Best	0.818	default
multiling-e5-large	Worst	0.562	Text pair similarity assessment:
	Best	0.848	default
	Worst	0.781	Measure semantic similarity between:

Table 14: Best and worst prompts for task: STS15

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.903	default
Qwen3-0.6B	Worst	0.821	Text similarity level in SemEval STS 2015 dataset:
	Best	0.914	default
bge-base-en-1.5	Worst	0.784	SemEval STS 2015 dataset - measure similarity between pairs:
	Best	0.88	default
bge-large-en-1.5	Worst	0.7	Measure semantic similarity in eng based on SemEval STS 2015 dataset:
	Best	0.88	default
bge-small-en-1.5	Worst	0.693	STS 2015 style semantic similarity evaluation:
	Best	0.873	default
multiling-e5-large	Worst	0.674	Measure text similarity in eng, SemEval STS 2015:
	Best	0.91	default
	Worst	0.855	Measure semantic similarity between:

Table 15: Best and worst prompts for task: STS22.v2

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.732	default
Qwen3-0.6B	Worst	0.685	Compare texts for semantic similarity:
	Best	0.718	default
bge-base-en-1.5	Worst	0.646	Query: assess text pair similarity
	Best	0.678	In eng, assess similarity of news article pair - SemEval 2022 Task 8:
bge-large-en-1.5	Worst	0.466	default
	Best	0.699	Similarity between texts in eng:
bge-small-en-1.5	Worst	0.489	default
	Best	0.674	Compare the similarity of these SemEval 2022 Task 8 news articles:
multiling-e5-large	Worst	0.469	default
	Best	0.69	default
	Worst	0.638	Similarity score for article pair:

Table 16: Best and worst prompts for task: StackExchangeClustering.v2

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.777	Cluster these StackExchange titles:
Qwen3-0.6B	Worst	0.556	Group similar texts:
	Best	0.712	default
bge-base-en-1.5	Worst	0.553	eng text for clustering of 25 sets, each with 10-50 classes and 100-1000 sentences per class:
	Best	0.58	Grouping for clustering:
bge-large-en-1.5	Worst	0.486	Encode for clustering: titles from 121 stackexchanges in eng
	Best	0.594	default
bge-small-en-1.5	Worst	0.5	Encode for clustering: titles from 121 stackexchanges in eng
	Best	0.573	default
multiling-e5-large	Worst	0.481	Group similar texts into clusters: Clustering of titles from 121 stackexchanges in eng
	Best	0.672	Cluster these StackExchange titles:
	Worst	0.6	default

Table 17: Best and worst prompts for task: Touche2020Retrieval.v3

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.644	Retrieval query:
Qwen3-0.6B	Worst	0.218	Representation for retrieval in eng, Touché Task 1: Argument Retrieval for Controversial Questions
	Best	0.699	default
bge-base-en-1.5	Worst	0.383	Encode (eng) for Touché Task 1: Argument Retrieval for Controversial Questions:
	Best	0.497	default
bge-large-en-1.5	Worst	0.035	Representation for retrieval in eng, Touché Task 1: Argument Retrieval for Controversial Questions
	Best	0.471	default
bge-small-en-1.5	Worst	0.014	Retrieve relevant argument in eng for Touché Task 1: Argument Retrieval for Controversial Questions:
	Best	0.576	default
multiling-e5-large	Worst	0.009	Retrieve relevant argument in eng for Touché Task 1: Argument Retrieval for Controversial Questions:
	Best	0.545	Search query:
	Worst	0.342	Encode for Touché Task 1 retrieval:

Table 18: Best and worst prompts for task: TweetSentimentClassification

Model	Type	Score	Prompt
KaLM-multiling-mini-2.5	Best	0.704	Classify the sentiment of this tweet:
Qwen3-0.6B	Worst	0.526	Category prediction:
	Best	0.706	Classify the sentiment of this tweet:
bge-base-en-1.5	Worst	0.483	Category prediction:
	Best	0.518	Label this example:
bge-large-en-1.5	Worst	0.417	default
	Best	0.529	Label the following:
bge-small-en-1.5	Worst	0.427	default
	Best	0.52	Classify this tweet in eng:
multiling-e5-large	Worst	0.417	default
	Best	0.645	eng sentiment analysis - classify:
	Worst	0.485	Determine the class for the given passage: