

GATech at AbjadMed: Bidirectional Encoders vs. Causal Decoders: Insights from 82-Class Arabic Medical Classification

Ahmed Khaled Khamis
Georgia Institute of Technology
akhamis6@gatech.edu

Abstract

This paper presents system description for Arabic medical text classification across 82 distinct categories. Our primary architecture utilizes a fine-tuned AraBERTv2 encoder enhanced with a hybrid pooling strategies, combining attention and mean representations, and multi-sample dropout for robust regularization. We systematically benchmark this approach against a suite of multilingual and Arabic-specific encoders, as well as several large-scale causal decoders, including zero-shot re-ranking via Llama 3.3 70B and feature extraction from Qwen 3B hidden states. Our findings demonstrate that specialized bidirectional encoders significantly outperform causal decoders in capturing the precise semantic boundaries required for fine-grained medical text classification. We show that causal decoders, optimized for next-token prediction, produce sequence-biased embeddings that are less effective for categorization compared to the global context captured by bidirectional attention. Despite significant class imbalance and label noise identified within the training data, our results highlight the superior semantic compression of fine-tuned encoders for specialized Arabic NLP tasks. Final performance metrics on the test set, including Accuracy and Macro-F1, are reported and discussed.

1 Introduction

AbjadMed shared task (Gupta et al., 2026) requires the classification of Arabic medical queries into 82 distinct categories. This high-cardinality classification task has a significant class imbalance, where training samples range from several hundred to as few as seven per category. Furthermore, manual inspection of the provided training data reveals a degree of label noise, where semantically similar categories (e.g., "General Medicine" vs. "Internal Medicine") are sometimes inconsistently assigned. These factors necessitate a system that is not only linguistically nuanced but also highly robust to variance and noise.

A central theme of current Natural Language Processing research is the tension between specialized, fine-tuned bidirectional encoders (such as BERT and its derivatives) (Vaswani et al., 2023) and large-scale causal decoders (such as the Llama/GPT family) (Touvron et al., 2023) (Brown et al., 2020). While the latter have demonstrated remarkable zero-shot reasoning capabilities, their "causal" nature—processing text in a single direction might be suboptimal for capturing the dense semantic boundaries required for 82-class medical categorization. In this paper, we investigate whether the massive parameters and generative pre-training of decoders can outperform task-specific fine-tuned encoders.

To address these challenges, we present our system for medical classification the AbjadMed shared task,¹ the system is designed for high-cardinality classification. Our primary approach leverages a fine-tuned AraBERTv2 encoder (Antoun et al., 2021). We enhance the standard classification architecture with hybrid pooling strategies, concatenating attention (Chen et al., 2022) and mean representations (Chen et al., 2018) to capture both global thematic context and specific medical keywords. To combat class imbalance and label noise, we implement Multi-Sample Dropout (Inoue, 2020) and Label Smoothing (Müller et al., 2020), providing a form of internal ensembling that stabilizes the decision boundaries for minority classes.

We contrast our primary system with several exploratory approaches, including a two-stage hybrid re-ranking pipeline using Llama 3.3 70B (Grattafiori et al., 2024) Instruct and feature extraction using Qwen 3 3B (Yang et al., 2025) embeddings. Our comparative analysis provides evidence that bidirectional encoders remain the superior choice for high-granularity Arabic medical

¹Code: <https://github.com/KickItLikeShika/abjadmed>

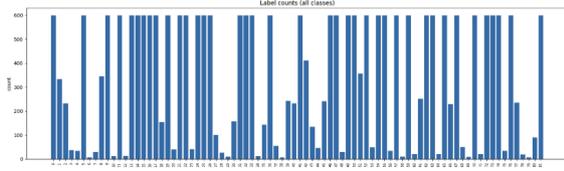


Figure 1: Training set class distribution

tasks, offering more effective semantic compression than their much larger causal counterparts.

2 Background

2.1 Dataset

The dataset provided for the shared task consists of Arabic medical queries and their corresponding category labels. The training set contains 27,951 samples distributed across 82 distinct medical categories. The test set comprises 18,634 unlabeled samples.

(Figure 1) shows the extreme class imbalance. While majority classes such as Hematological diseases, Addiction, and Neurological diseases are represented by 600 samples each, minority classes such as Biochemistry, Vascular surgery, and In vitro fertilization (IVF) contain as few as 7 samples. This long-tail distribution poses a significant challenge for standard cross-entropy loss, as the model may easily overfit to majority class features while failing to generalize for rarer specialties.

2.2 Label Noise and Ambiguity

Beyond the quantitative imbalance, a qualitative analysis of the training data revealed challenges regarding label consistency. During our exploratory phase, we identified several instances of label noise, where the assigned gold-standard label appeared inconsistent with the textual content of the query. Examples include: 1) Queries clearly describing dermatological symptoms were occasionally labeled as General Medicine rather than Dermatological diseases, 2) Mismatch between categories like Sexual Health and Medicinal herbs. This inherent noise suggests that the dataset contains "soft" boundaries between certain categories. Consequently, we opted for Label Smoothing during training to prevent the model from becoming overly confident in potentially incorrect labels, thereby encouraging better generalization across these ambiguous semantic regions.

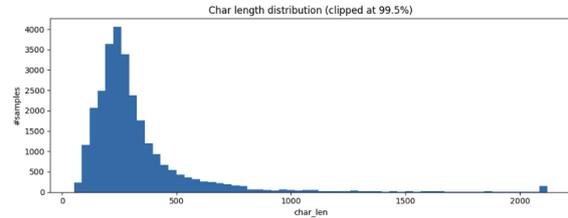
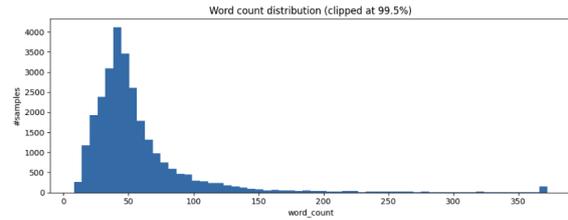


Figure 2: Word and character count distribution clipped at 99.5%

2.3 Data Preprocessing

Given the varied length of medical queries (Figure 2), we utilized a dynamic padding strategy (Toker et al., 2025). Rather than padding all sequences to a fixed global maximum length, we padded each batch to the length of its longest sequence. This reduced the computational overhead of processing unnecessary padding tokens, allowing for faster iterations and more efficient GPU utilization during the fine-tuning process.

3 System Architecture

Our system follows a fine-tuning paradigm where a pre-trained bidirectional encoder is augmented with a specialized classification head and hybrid pooling layers.

3.1 Encoder Selection

After benchmarking multiple architectures, we selected AraBERTv2 (bert-base-arabertv02) as our primary encoder. AraBERT is pre-trained on a massive corpus of Arabic text, making it particularly effective at capturing the semantic nuances of medical Modern Standard Arabic (MSA). We utilize the transformer’s hidden states as the foundation for our discriminative features.

3.2 Hybrid Pooling Strategy

Standard text classification often relies solely on the representation of the [CLS] token (Wang et al., 2024). However, for medical queries where diagnostic evidence may be scattered across the sequence, a more comprehensive aggregation is required. We implemented a Hybrid Pooling

layer that concatenates two distinct representations: **Mean Pooling:** Calculates the average of all token embeddings (excluding padding), providing a global thematic summary of the query. **Attention Pooling:** A learnable attention mechanism that assigns importance scores to individual tokens. This allows the model to "focus" on salient medical keywords (e.g., specific symptoms or organ names) while down-weighting irrelevant stop words. The resulting feature vector $h_{pooled} = [h_{mean}; h_{attn}]$ doubles the hidden dimensionality ($hidden \times 2$), providing a richer signal to the classification head.

3.3 Classification Head

The concatenated features are passed through a dense layer with Layer Normalization (Ba et al., 2016) and GELU (Hendrycks and Gimpel, 2023) activation. To address the dual challenges of class imbalance and label noise, we employ Multi-Sample Dropout. Unlike standard dropout, which applies a single mask, Multi-Sample Dropout creates five parallel dropout paths with varying rates (0.1 to 0.3) during training. The final classification layer is applied to each of these paths, and the resulting logits are averaged. This technique acts as an internal ensemble, effectively regularizing the model and reducing the variance of predictions, which is critical for the stability of minority class boundaries.

3.4 Training Methodology

To preserve the linguistic knowledge of the pre-trained encoder while allowing the classification head to adapt to the 82-class schema, we utilized several advanced optimization techniques: **Layer-wise Learning Rate Decay (LLRD):** We applied a decay factor of 0.95, ensuring that layers closer to the input (embeddings) adapt more slowly than the task-specific top layers. **Label Smoothing:** We applied a smoothing factor of 0.1 to the cross-entropy loss. This prevents the model from overfitting to the "hard" labels, which we identified as potentially noisy, and encourages the learning of more flexible decision boundaries. **AdamW Optimizer:** Coupled with a Cosine Learning Rate Scheduler, starting at a base learning rate of 2×10^{-5} (Loshchilov and Hutter, 2019).

4 Experimental Setups

4.1 Encoders

We conducted a systematic evaluation of several Transformer-based encoders, spanning both Arabic-specific and multilingual models.

AraBERTv2 (bert-base-arabertv02) consistently emerged as the top performer. Its specialized pre-training on Arabic-specific corpora allows it to better capture the linguistic nuances of medical Modern Standard Arabic (MSA) compared to multilingual alternatives. multilingual-E5-large (?) provided the second-best performance. We attribute its strength to its contrastive pre-training on large-scale text pairs, which yields robust semantic representations across languages.

CamelBERT (Inoue et al., 2021) and mDeBERTa-v3 (He et al., 2021) showed competitive performance but were slightly less effective at capturing the specific medical vocabulary of the shared task. EuroBERT (Boizard et al., 2025) and mmBERT (Marone et al., 2025) lagged significantly behind, likely due to a lack of representative Arabic medical data in their initial pre-training sets.

4.2 The Semantic Limitations of Causal Decoders

We investigated whether the massive parameter count and extensive pre-training of the Qwen 3 3B model could provide superior semantic embeddings compared to smaller BERT-style encoders. We extracted the hidden states from the final layer of the model to serve as static features for a task-specific classification head.

Our results showed that this approach significantly underperformed compared to the fine-tuned AraBERT. We attribute this to the fundamental difference between Causal and Bidirectional Attention: **Sequence Bias:** Causal decoders are optimized for next-token prediction, which biases their internal representations towards the sequence history. **Semantic Compression:** Unlike bidirectional encoders, which allow every token to attend to the entire context simultaneously, the decoder's hidden states are "generative" in nature. They fail to compress the full semantic meaning of a medical query into a dense vector suitable for discriminative 82-class categorization. Fine-tuned encoders produce sharper semantic boundaries that are essential for distinguishing between similar medical sub-specialties.

4.3 Zero-shot Hybrid Re-ranking

Finally, we explored a two-stage hybrid pipeline. In this setup, AraBERT was used as a candidate generator to propose the Top-15 most probable categories. These candidates, along with the original text, were passed to Llama 3.3 70B Instruct via a structured zero-shot prompt, asking the model to make the final selection. This approach introduced more noise than it resolved. The LLM frequently encountered a "Schema Mismatch": while its choices were often semantically logical (e.g., choosing Dermatology for a skin-related query), they often conflicted with the specific boundaries of the shared task's 82-class schema (where the correct label might be Skin and Beauty). While the LLM possesses superior general reasoning, the fine-tuned AraBERT encoder more accurately learned the specific intent and labeling conventions of the task's annotators.

5 Results

We evaluated the performance of our system on the official competition public test set. Given the extreme class imbalance across the 82 categories, we utilize the Macro-F1 score as our primary performance metric, as it provides an unweighted average that accounts for the model's ability to generalize to minority classes. Table 1 summarizes the performance of our primary architecture compared to our exploratory benchmarks and the hybrid re-ranking pipeline. Our fine-tuned AraBERTv2 configuration, incorporating hybrid pooling and multi-sample dropout, achieved the highest performance. Interestingly, the addition of a large-scale decoder for re-ranking (Llama 3.3 70B) resulted in a performance degradation, highlighting the "schema mismatch" between general reasoning models and task-specific classification boundaries.

Model Configuration	Macro-F1
AraBERTv2	0.3934
multilingual-E5-large	0.3804
AraBERTv2+Llama 3.3 70B	0.3035
CamelBert	0.3603
Qwen 3 3B (Feature Extraction)	0.1278

Table 1: System performance on the official test set. The primary AraBERTv2 configuration utilizes our proposed hybrid pooling and multi-sample dropout head.

6 Discussion and Limitations

6.1 The Encoder Advantage in High-Cardinality Tasks

Our results reinforce the hypothesis that for high-cardinality classification within a specific domain, discriminative fine-tuning on a bidirectional encoder is more critical than the general reasoning capabilities of large causal decoders. While models such as Llama 3.3 70B contain extensive medical knowledge, they are not aligned with the fine-grained annotation boundaries of this shared task. In contrast, fine-tuning allows AraBERT to sharpen semantic decision boundaries and distinguish between medically similar but distinctly labeled categories (e.g., Alternative medicine vs. Medicinal herbs).

6.2 The Inefficacy of Generative Embeddings

A key finding of this study is the poor performance of feature extraction from causal decoders such as Qwen 3 3B and Llama 3 8B. We observe that causal attention can be suboptimal for sequence-level semantic compression, as the final token embedding often reflects recent sequence history rather than the full input. Bidirectional encoders, through all-to-all attention, allow medical keywords appearing anywhere in the query to contribute equally to the final representation, which is especially important when diagnostic cues occur early in long, conversational medical queries.

6.3 Future Work

While our modeling experiments combine several training and architectural techniques, a more fine-grained analysis remains valuable. Future work should include ablation studies that evaluate the individual impact of hybrid pooling, multi-sample dropout, and label smoothing by introducing each component independently. Such controlled ablations would provide clearer insight into the contribution of each technique.

7 Conclusion

We presented a system for medical text classification in the AbjadNLP shared task, addressing the challenges of extreme class imbalance across an 82-class label space. Our approach, based on fine-tuning an AraBERTv2 encoder with hybrid pooling and multi-sample dropout, achieved strong performance compared to multilingual encoders and large causal decoders.

References

- Wissam Antoun, Fady Baly, and Hazem Hajj. 2021. [Arabert: Transformer-based model for arabic language understanding](#). *Preprint*, arXiv:2003.00104.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *Preprint*, arXiv:1607.06450.
- Nicolas Boizard, Hippolyte Gisserot-Boukhlef, Duarte M. Alves, André Martins, Ayoub Hammal, Caio Corro, Céline Hudelot, Emmanuel Malherbe, Etienne Malboeuf, Fanny Jourdan, Gabriel Hautreux, João Alves, Kevin El-Haddad, Manuel Faysse, Maxime Peyrard, Nuno M. Guerreiro, Patrick Fernandes, Ricardo Rei, and Pierre Colombo. 2025. [Eurobert: Scaling multilingual encoders for european languages](#). *Preprint*, arXiv:2503.05500.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Fang Chen, Gourav Datta, Souvik Kundu, and Peter Beerel. 2022. [Self-attentive pooling for efficient deep learning](#). *Preprint*, arXiv:2209.07659.
- Qian Chen, Zhen-Hua Ling, and Xiaodan Zhu. 2018. [Enhancing sentence embedding with generalized pooling](#). *Preprint*, arXiv:1806.09828.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Pranav Gupta, Niranjan Kumar M, Balaji Nagarajan, Imed Zitouni, and Mo El-Haj. 2026. [Abjadmed: Arabic medical text classification at abjadnlp 2026](#). In *Proceedings of the 2nd Workshop on NLP for Languages Using Arabic Script (AbjadNLP 2026)*, co-located with the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026), Rabat, Morocco.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced bert with disentangled attention](#). *Preprint*, arXiv:2006.03654.
- Dan Hendrycks and Kevin Gimpel. 2023. [Gaussian error linear units \(gelus\)](#). *Preprint*, arXiv:1606.08415.
- Go Inoue, Bashar Alhafni, Nurpeiis Baimukan, Houda Bouamor, and Nizar Habash. 2021. The interplay of variant, size, and task type in Arabic pre-trained language models. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, Kyiv, Ukraine (Online). Association for Computational Linguistics.
- Hiroshi Inoue. 2020. [Multi-sample dropout for accelerated training and better generalization](#). *Preprint*, arXiv:1905.09788.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). *Preprint*, arXiv:1711.05101.
- Marc Marone, Orion Weller, William Fleshman, Eugene Yang, Dawn Lawrie, and Benjamin Van Durme. 2025. [mmbert: A modern multilingual encoder with annealed language learning](#). *Preprint*, arXiv:2509.06888.
- Rafael Müller, Simon Kornblith, and Geoffrey Hinton. 2020. [When does label smoothing help?](#) *Preprint*, arXiv:1906.02629.
- Michael Toker, Ido Galil, Hadas Orgad, Rinon Gal, Yoad Tewel, Gal Chechik, and Yonatan Belinkov. 2025. [Padding tone: A mechanistic analysis of padding tokens in t2i models](#). *Preprint*, arXiv:2501.06751.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2023. [Attention is all you need](#). *Preprint*, arXiv:1706.03762.
- Ao Wang, Fengyuan Sun, Hui Chen, Zijia Lin, Jungong Han, and Guiguang Ding. 2024. [\[cls\] token tells everything needed for training-free efficient mllms](#). *Preprint*, arXiv:2412.05819.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.