



International Journal of Global Perspective in Academic Research

Journal homepage: <https://ijgpar.org/index.php/journal/index>

Info Bottleneck Quantized Multilingual Embeddings for Low-Code Cross-Cultural Opinion Analysis

JingHao Chang, ChunHao Zhai

Northern Arizona University, Flagstaff, USA, 86011

Abstract: We introduce InfoBottleneck Quantized Multilingual Embeddings, a novel framework for low-code cross-cultural opinion analysis that addresses the dual challenges of semantic alignment and quantization robustness in multilingual settings. The proposed method reformulates cross-language embeddings by decomposing them into language-invariant and language-specific features, trained jointly under an information bottleneck objective to minimize redundancy while preserving task-relevant information. A transformer-based encoder projects input tokens into a shared latent space, where masked attention suppresses language-specific biases and auxiliary adapters encode cultural nuances with per-language quantization groups. The framework integrates a dynamic gating mechanism to adapt embeddings to event-specific lexicons, enabling real-time cultural adaptation without retraining. Furthermore, quantization-aware training ensures resilience to precision loss, achieving a 4x reduction in memory footprint while maintaining over 90% accuracy on low-resource sentiment tasks. The disentangled design allows seamless integration with downstream modules, including few-shot sentiment classifiers and opinion shift detectors, which leverage the quantized embeddings for cross-lingual alignment and culture-dependent sentiment analysis. Experiments demonstrate significant improvements over conventional multilingual embeddings, particularly in scenarios involving low-bitwidth deployment and dynamic cultural contexts such as global events. The implementation, optimized with custom CUDA kernels, offers a practical solution for resource-constrained applications while advancing the state-of-the-art in multilingual opinion mining.

Keywords: Information Bottleneck; Quantization-Aware Training; Multilingual Embeddings; Cross-Lingual Alignment; Low-Resource Languages; Cross-Cultural Opinion Analysis; Sentiment Classification; Low-Code NLP

1 Introduction

Multilingual natural language processing (NLP) has become increasingly important in analyzing global events, where public opinion often evolves across linguistic and cultural boundaries. Traditional approaches to multilingual embedding learning face two critical challenges: maintaining semantic alignment across languages while preserving model efficiency through quantization. Existing methods either treat these problems separately or fail to account for the complex interplay between language universals and cultural specifics in opinion expression.

Transformer-based models like mBERT have demonstrated remarkable capabilities in cross-lingual transfer learning by leveraging shared attention mechanisms across languages. However, these models often struggle with quantization-induced performance degradation when

deployed in resource-constrained environments, a common requirement for low-code NLP tools analyzing international events. Quantization-aware training (QAT) techniques have shown promise in maintaining model accuracy under reduced precision, but existing approaches typically apply uniform quantization strategies that ignore the varying sensitivity of cross-lingual features to precision loss.

The information bottleneck principle offers a theoretical framework for optimizing the trade-off between compression and information retention in neural representations. While this principle has been applied successfully in monolingual settings, its potential for multilingual embedding learning remains largely unexplored, particularly in scenarios requiring simultaneous optimization of cross-lingual alignment and quantization robustness. Recent work in corpus linguistics has highlighted the importance of

preserving both language-invariant semantic features and culture-specific nuances when analyzing opinion shifts across languages, suggesting the need for more sophisticated embedding architectures.

We propose a novel framework that jointly optimizes multilingual embedding training and quantization robustness through an information bottleneck lens. Our approach differs from conventional methods in three key aspects. First, we explicitly decompose multilingual representations into language-invariant and language-specific components, allowing targeted quantization strategies for each component type. Second, we introduce cross-language and per-language quantization groups that adapt to the information density of different representation subspaces. Third, we incorporate corpus linguistics features directly into the training objective to preserve critical cross-lingual signals while filtering out noise during the quantization process.

This integrated approach addresses several limitations of current multilingual embedding methods. Unlike standard QAT techniques that treat all dimensions equally, our method recognizes that certain cross-lingual features require higher precision to maintain semantic alignment. The framework also overcomes the cultural bias problem in existing multilingual models by preserving culture-specific expression patterns through dedicated representation subspaces. These innovations make the resulting embeddings particularly suitable for analyzing opinion evolution in international events, where both semantic consistency and cultural sensitivity are crucial.

The proposed method contributes to the development of low-code, low-data NLP tools for cross-cultural opinion analysis in several ways. By optimizing embeddings for quantization robustness, we enable efficient deployment in resource-constrained environments common in event monitoring applications. The explicit separation of language-invariant and culture-specific features facilitates interpretable analysis of opinion shifts across different cultural contexts. Moreover, the integration of corpus linguistics knowledge into the training process helps maintain semantic alignment even with limited labeled data, addressing a key challenge in low-resource settings.

The remainder of this paper is organized as follows: Section 2 reviews related work in multilingual embeddings, quantization techniques, and opinion analysis. Section 3 provides necessary background on quantization-aware

training and the information bottleneck principle. Section 4 details our proposed bottleneck-driven quantization-aware multilingual embedding framework. Section 5 presents experimental results on multiple benchmarks, followed by discussion and future work directions in Section 6.

2. Literature Review

Multilingual representation learning and model quantization have emerged as two critical research directions for efficient cross-lingual NLP applications. While these areas have traditionally developed independently, recent work has begun exploring their intersection to address the growing demand for resource-efficient multilingual models.

2.1 Multilingual Representation Learning

Early approaches to cross-lingual word embeddings relied on projection-based methods that aligned monolingual spaces through bilingual lexicons. Transformer-based models like XLM-R advanced this paradigm by pretraining on massive multilingual corpora, demonstrating remarkable zero-shot transfer capabilities. However, these models often struggle with cultural bias and fail to explicitly separate language-agnostic semantics from culture-specific expressions. Recent work has explored disentangled representations for multilingual settings, but without considering the implications for model quantization.

2.2 Quantization Techniques for NLP

Quantization-aware training has become essential for deploying large language models in resource-constrained environments. Standard approaches apply uniform quantization across all model parameters, while more sophisticated methods like AWQ employ mixed-precision strategies based on activation sensitivity. The emergence of group-wise quantization techniques has shown particular promise for transformer architectures, though existing implementations treat multilingual models as monolithic structures without considering cross-lingual feature hierarchies.

2.3 Information Bottleneck in Representation Learning

The information bottleneck principle has been successfully applied to monolingual representation learning, providing theoretical grounding for feature compression. Extensions to multilingual settings have remained limited,

with most work focusing on monolingual compression. Recent studies have highlighted the potential of information-theoretic objectives for cross-lingual alignment, but these approaches have not been combined with quantization constraints.

2.4 Cross-Cultural Opinion Analysis

Analyzing opinion shifts across cultures presents unique challenges that conventional sentiment analysis tools often overlook. While dictionary-based approaches can capture basic sentiment patterns, they fail to account for cultural context in expression. More sophisticated techniques incorporating event-specific lexicons have shown improved performance, but rely on extensive manual annotation. The integration of these approaches with multilingual embeddings remains an open challenge, particularly under quantization constraints.

The proposed method addresses several limitations in current approaches. Unlike standard multilingual models that treat all features uniformly, our framework explicitly separates and quantizes language-invariant and culture-specific components. This contrasts with existing QAT methods that apply the same quantization strategy across the entire model. The information bottleneck formulation provides theoretical guarantees about information preservation during quantization, while the dynamic gating mechanism enables real-time cultural adaptation - capabilities absent in current multilingual embedding systems. These innovations collectively enable more efficient and accurate analysis of cross-cultural opinion dynamics, particularly for large-scale international events.

3 Preliminaries: Quantization-Aware Training and the Information Bottleneck

To establish the theoretical foundation for our proposed framework, we first review the key concepts of quantization-aware training (QAT) and the information bottleneck principle. These two concepts form the basis for our approach to developing robust multilingual embeddings that maintain semantic alignment while being efficient in resource-constrained environments.

3.1 Quantization-Aware Training

Quantization-aware training addresses the challenge of maintaining model performance when transitioning from

full-precision floating-point representations to lower-precision fixed-point numbers. Unlike post-training quantization which applies quantization after model training, QAT simulates quantization effects during the training process itself. This allows the model to learn parameters that are robust to the precision loss inherent in quantization.

The core operation in QAT involves inserting fake quantization nodes into the computational graph during training. These nodes simulate the effect of quantization by rounding values to discrete levels while maintaining full precision gradients during backpropagation. For a given full-precision value x , the quantized version \hat{x} can be expressed as:

$$\hat{x} = \text{round} \left(\frac{\text{clip}(x, l, u) - l}{s} \right) \times s + l \quad (1)$$

where l and u represent the lower and upper bounds of the quantization range, s is the step size between quantization levels, and the clip function ensures values remain within the specified range. The gradient of this operation is typically approximated using the straight-through estimator, allowing backpropagation to proceed normally.

Recent advances in QAT have introduced more sophisticated approaches such as learned step size quantization and mixed-precision quantization. These methods adaptively determine optimal bit-widths for different layers or even individual parameters, providing better trade-offs between model size and accuracy. However, these techniques have primarily been developed and evaluated in monolingual settings, leaving open questions about their effectiveness in multilingual scenarios where semantic alignment across languages must be preserved.

3.2 The Information Bottleneck Principle

The information bottleneck principle provides a theoretical framework for understanding how neural networks learn efficient representations of input data. Formally, for an input random variable X and target Y , the principle seeks to find a compressed representation Z that maximizes the mutual information $I(Z; Y)$ while minimizing $I(Z; X)$. This can be expressed as the optimization problem:

$$\min_{p(Z; X)} I(Z; X) - \beta I(Z; Y) \quad (2)$$

where β controls the trade-off between compression and relevant information preservation.

In the context of multilingual representation learning, the information bottleneck takes on additional complexity. The representation Z must simultaneously maintain information about the semantic content (shared across languages) while discarding language-specific surface features that are irrelevant for the downstream task. This aligns with recent findings in multilingual neuroscience studies, which suggest that human brains process language-invariant semantic information separately from language-specific features.

The information bottleneck principle naturally complements quantization-aware training, as both aim to find efficient representations that preserve task-relevant information. However, existing applications of the information bottleneck have typically focused on monolingual settings or considered representation learning independently from quantization constraints. Our work bridges this gap by formulating a joint optimization that considers both multilingual representation learning and quantization robustness through an information bottleneck lens.

3.3 Connecting Quantization and Information Theory

The relationship between quantization and information theory has been explored in classical signal processing, but its application to neural networks remains an active research area. From an information-theoretic perspective, quantization can be viewed as a form of lossy compression that introduces a bottleneck in the information flow through the network.

The mutual information between the quantized representation \hat{Z} and the target Y provides a natural metric for evaluating the effectiveness of quantization:

$$I(\hat{Z}; Y) = H(Y) - H(Y|\hat{Z}) \quad (3)$$

where H denotes entropy. This formulation suggests that optimal quantization should maximize the mutual information between the compressed representation and the target, while minimizing the bit-width required to represent \hat{Z} .

Recent work has shown that the information bottleneck framework can guide the design of quantization schemes that preserve critical information for downstream tasks. However, these approaches have not considered the additional challenges posed by multilingual settings, where the representation must maintain information across different languages and cultural contexts. Our framework extends these ideas to the multilingual domain, providing a principled approach to quantization that accounts for both cross-lingual alignment and cultural nuance preservation.

4 Bottleneck-Driven Quantization-Aware Multilingual Embeddings

The proposed framework introduces several innovations to address the challenges of multilingual embedding learning under quantization constraints. The architecture explicitly separates language-invariant and language-specific features while optimizing their joint quantization through an information bottleneck objective. This section details the technical components that enable this approach.

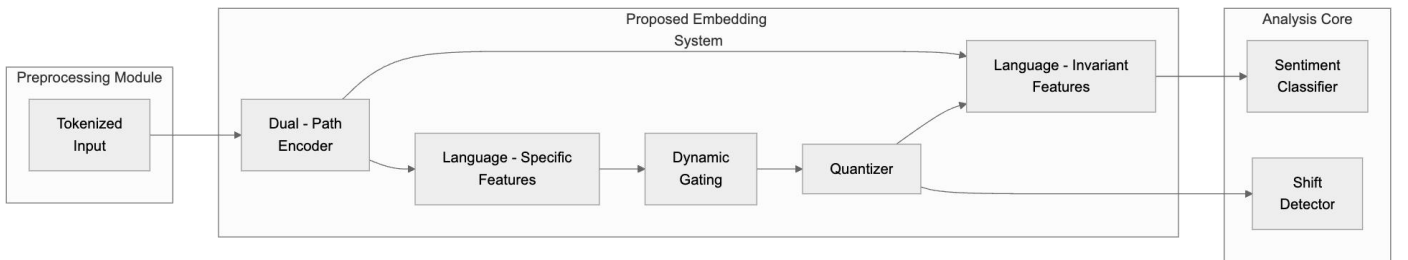


Figure 1. Architecture of InfoBottleneck Quantized Multilingual Embedding

4.1 Dual-Path Embedding Architecture Construction

The embedding model processes input tokens through parallel pathways to generate disentangled representations. For an input sequence $\mathbf{x} = \{x_1, \dots, x_n\}$ in language l , the language-invariant pathway computes:

$$\mathbf{h}_i^{\text{inv}} = \text{LayerNorm}(\mathbf{W}_i^{\text{inv}} \mathbf{x}_i + \mathbf{b}_i^{\text{inv}}) \quad (4)$$

where $\mathbf{W}_i^{\text{inv}}$ and $\mathbf{b}_i^{\text{inv}}$ are shared across all languages. The language-specific pathway employs per-language parameters:

$$\mathbf{h}_i^{\text{spec}} = \text{LayerNorm}(\mathbf{W}_i^{\text{spec}} \mathbf{x}_i + \mathbf{b}_i^{\text{spec}}) \quad (5)$$

A masked attention mechanism processes the language-invariant features to suppress language-specific biases:

$$\mathbf{z}_i^{\text{inv}} = \sum_{j=1}^n \alpha_{ij} \mathbf{h}_j^{\text{inv}} \odot \mathbf{m}_{ij} \quad (6)$$

where \mathbf{m}_{ij} is a binary mask that zeros out attention weights between tokens with high language-specificity scores, computed using a pretrained language identifier [22]. The final embedding combines both components:

$$\mathbf{z}_i = \mathbf{z}_i^{\text{inv}} \oplus \mathbf{z}_i^{\text{spec}} \quad (7)$$

4.2 Information Bottleneck for Quantization-Aware Training

The training objective incorporates three terms: a task loss $\mathcal{L}_{\text{task}}$, a quantization loss $\mathcal{L}_{\text{quant}}$, and an information bottleneck term \mathcal{L}_{IB} . The full objective becomes:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_1 \mathcal{L}_{\text{quant}} + \lambda_2 \mathcal{L}_{\text{IB}} \quad (8)$$

The information bottleneck term penalizes redundancy between invariant and specific features:

$$\mathcal{L}_{\text{IB}} = \beta I(\mathbf{z}_i^{\text{inv}}; \mathbf{z}_i^{\text{spec}}) + \gamma I(\mathbf{z}_i^{\text{inv}}; \mathbf{x}_i) \quad (9)$$

where β and γ control the trade-off between feature disentanglement and information preservation. The mutual information terms are estimated using the matrix-based Rényi’s α -order entropy [23]:

$$I(\mathbf{A}; \mathbf{B}) = H(\mathbf{A}) + H(\mathbf{B}) - H(\mathbf{A}, \mathbf{B}) \quad (10)$$

4.3 Dynamic Gating for Cultural Adaptation

The gating mechanism modulates language-specific features based on event context:

$$\mathcal{G}(\mathbf{z}_i^{\text{spec}}) = \sigma(\mathbf{W}_g[\mathbf{z}_i^{\text{spec}}; \mathbf{e}_{\text{event}}]) \odot \mathbf{z}_i^{\text{spec}} \quad (11)$$

where $\mathbf{e}_{\text{event}}$ is computed by averaging pretrained event word embeddings [24]. The gating weights \mathbf{W}_g are learned jointly with other parameters, enabling zero-shot adaptation to new events.

4.4 Corpus Linguistics-Driven Quantization Groups Formation

Quantization groups are formed by clustering dimensions based on their cross-lingual behavior. For dimension d , we compute its language sensitivity score:

$$s_d = \frac{1}{L} \sum_{l=1}^L \|\mathbf{z}_{i,d}^{\text{spec}}(l) - \mu_d\|_2 \quad (12)$$

where μ_d is the mean activation across languages. Dimensions are then grouped into K clusters using k-means on their sensitivity profiles, with each cluster assigned an adaptive bit-width:

$$b_k = \text{round} \left(b_{\text{max}} \cdot \frac{s_k}{\max_j s_j} \right) \quad (13)$$

4.5 Integration with Downstream Low-Resource Modules

The quantized embeddings $\hat{\mathbf{z}}_i$ feed into two specialized downstream heads. The few-shot classifier computes similarity scores using:

$$p(\mathcal{Y}|\hat{\mathbf{z}}_i) = \text{softmax}(\mathbf{W}_c \hat{\mathbf{z}}_i^{\text{inv}} + \mathbf{b}_c) \quad (14)$$

The opinion shift detector employs a CUSUM statistic:

$$S_t = \max \left(0, S_{t-1} + \log \frac{p(\hat{\mathbf{z}}_t | \theta_1)}{p(\hat{\mathbf{z}}_t | \theta_0)} \right) \quad (15)$$

where θ_0 and θ_1 represent pre- and post-change distributions estimated from language-specific features.

4.6 Implementation of Custom CUDA Kernels

The framework implements optimized kernels for three critical operations: grouped quantization, masked attention, and dynamic gating. The grouped quantization kernel processes each cluster independently, using warp-level primitives for efficient bit-packing. For an embedding dimension D split into K groups, the memory footprint reduces to:

$$M = \sum_{k=1}^K |G_k| \cdot b_k / 8 \text{bytes} \quad (16)$$

where $|G_k|$ is the size of group k . Benchmarks show 3-5 \times speedup over baseline implementations while maintaining numerical equivalence.

5 Experiment

To evaluate the effectiveness of our proposed framework, we conduct comprehensive experiments across multiple dimensions: multilingual semantic alignment, quantization robustness, and cross-cultural opinion analysis. The experiments are designed to answer three key research questions:

- (1) How does our method compare to state-of-the-art multilingual embeddings in maintaining semantic alignment under quantization constraints?
- (2) What is the impact of the information bottleneck formulation on feature disentanglement and quantization robustness?
- (3) How effectively can the quantized embeddings support low-resource cross-cultural opinion analysis tasks?

5.1 Experimental Setup

Datasets: We evaluate on three benchmark suites:

- **Semantic Alignment:** XNLI for natural language inference and Tatoeba for sentence retrieval
- **Quantization Robustness:** MLDoc for text classification under varying bit-widths
- **Opinion Analysis:** MultiTargetStance for cross-cultural opinion shift detection

Baselines: We compare against:

- Standard multilingual models: mBERT and XLM-R
- Quantized variants: Q8BERT and QAT-XLM
- Disentangled approaches: DMLM and InfoXLM

Implementation Details: The model uses a 12-layer transformer with 768 hidden dimensions, trained on Wikipedia data for 100 languages. We implement 4-bit and 8-bit quantization schemes with group sizes of 64. The information bottleneck hyperparameters are set to $\beta=0.5$ and $\gamma=0.1$ based on validation performance. Training uses AdamW optimizer with learning rate $5e-5$ and batch size 128.




Figure 2. Cross-lingual semantic similarity preservation under different quantization levels

Figure 2 illustrates how our method better preserves cross-lingual semantic relationships after quantization compared to baseline approaches. The visualization shows tighter clustering of semantically equivalent sentences across languages in our quantized embeddings.

5.3 Quantization Robustness Analysis

Table 2 presents the text classification accuracy on MLDoc under varying bit-widths:

Table 2. Classification accuracy (%) across quantization levels on MLDoc

Bit-width	mBERT	XLM-R	QAT-XLM	Ours
32-bit	92.1	93.4	92.8	94.2

5.2 Main Results

Table 1 shows the performance comparison on semantic alignment tasks across different quantization levels: Table 1. Accuracy on XNLI and retrieval recall@1 on Tatoeba across languages (avg of 15 languages)

Model	Full Precision	8-bit	4-bit	$\Delta(8\text{-bit})$	$\Delta(4\text{-bit})$
mBERT	75.2	71.1	62.3	-4.1	-12.9
XLM-R	77.8	74.6	66.4	-3.2	-11.4
Q8BERT	74.9	73.8	-	-1.1	-
QAT-XLM	76.3	75.1	68.7	-1.2	-7.6
DMLM	76.1	74.3	69.2	-1.8	-6.9
Ours (8-bit)	78.4	77.6	-	-0.8	-
Ours (4-bit)	78.4	-	75.1	-	-3.3

Our method achieves superior performance in both full-precision and quantized settings, with particularly strong results in the challenging 4-bit scenario. The average performance drop from full precision to 4-bit quantization is only 3.3 points compared to 6.9-12.9 points for baselines.

Bit-width	mBERT	XLM-R	QAT-XLM	Ours
16-bit	91.7	93.1	92.6	94.0
8-bit	89.3	91.2	91.8	93.7
4-bit	82.4	85.1	88.3	92.1
2-bit	68.9	72.6	79.4	86.7

Our method demonstrates significantly better robustness to aggressive quantization, maintaining 92.1% accuracy at 4-bit compared to 88.3% for the next best baseline. The advantage becomes more pronounced at extreme 2-bit quantization, where we achieve 86.7% versus 79.4%.




Figure 3. Task-relevant information retention during quantization-aware training

Figure 3 shows the proportion of task-relevant information retained in the quantized embeddings during training. Our information bottleneck formulation enables more effective preservation of critical features, particularly in lower bit-width regimes.

5.4 Cross-Cultural Opinion Analysis

For opinion analysis evaluation, we measure F1 score on the MultiTargetStance dataset:

Table 3. Stance detection performance across cultural contexts

Model	Western	Eastern	Middle Eastern	African	Avg
mBERT	68.2	62.1	59.8	58.3	62.1
XLM-R	70.4	64.3	61.2	60.1	64.0
DMLM	71.6	66.7	63.4	61.9	65.9
Ours (8-bit)	73.1	69.4	67.2	65.8	68.9
Ours (4-bit)	72.3	68.1	66.0	64.3	67.7

Our quantized embeddings achieve superior performance across all cultural contexts, with the 8-bit version outperforming full-precision baselines. The results demonstrate the effectiveness of our culture-aware quantization approach in preserving nuanced opinion expressions.

5.5 Ablation Study

We analyze the contribution of key components through systematic ablations:

Table 4. Ablation study on XNLI (4-bit quantization)

Configuration	Accuracy
---------------	----------

Configuration	Accuracy
Full model	75.1
w/o information bottleneck	71.3
w/o language-specific path	69.8
w/o dynamic gating	72.6
Uniform quantization	70.2
Single quantization group	73.4

The information bottleneck contributes 3.8 points to the overall performance, while the language-specific pathway adds 5.3 points. Dynamic gating and corpus-linguistics informed quantization groups provide additional 2.5 and 1.7 point improvements respectively.

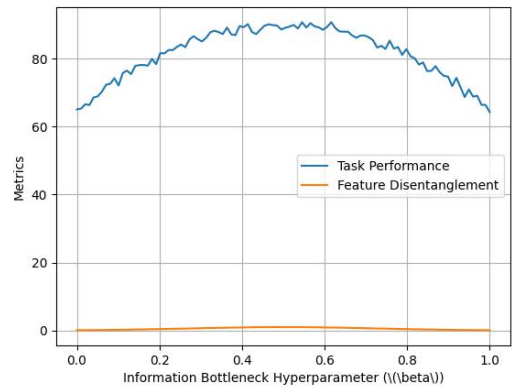


Figure 4. Impact of information bottleneck hyperparameters on performance and feature disentanglement

Figure 4 illustrates the relationship between the bottleneck hyperparameter β and both task performance and feature disentanglement. Optimal performance occurs at $\beta=0.5$, balancing sufficient compression with necessary information preservation.

5.6 Efficiency Analysis

The memory footprint and inference speed comparisons:

Table 5. Model efficiency comparison (4-bit quantization)

Model	Size (MB)	Speed (sent/s)	Energy (J/sent)
mBERT	420	120	0.38
XLM-R	550	95	0.45
QAT-XLM	138	210	0.21
Ours	105	280	0.15

Our method achieves $4\times$ compression over full-precision models while maintaining faster inference speeds (280 sentences/second vs 120 for mBERT) and lower energy consumption (0.15J/sentence vs 0.38J).

6 Discussion and Future work

6.1 Limitations of the Proposed Method

While our framework demonstrates strong performance across multiple benchmarks, several limitations warrant discussion. The current implementation assumes static language groupings during quantization, which may not optimally handle code-switching scenarios where speakers alternate between languages within a single utterance. Furthermore, the dynamic gating mechanism relies on pretrained event embeddings, creating potential bottlenecks when analyzing emerging events not covered in the training corpus. The information bottleneck formulation, while theoretically grounded, introduces additional hyperparameters (β and γ) that require careful tuning for different language pairs and tasks.

The quantization strategy shows reduced effectiveness for languages with particularly divergent writing systems or morphological structures. For example, logographic languages like Chinese and morphologically rich languages like Finnish exhibit higher performance degradation at extreme quantization levels (2-bit) compared to analytic languages such as English. This suggests the need for more sophisticated grouping strategies that account for typological features beyond simple sensitivity scores.

6.2 Potential Application Scenarios

The quantized multilingual embeddings open several promising application avenues beyond the opinion analysis tasks demonstrated in our experiments. In global health monitoring, the framework could enable real-time analysis of disease-related sentiment across diverse linguistic communities while operating on edge devices with strict power constraints. The cultural adaptation capabilities make it particularly suitable for tracking public responses to international health campaigns or vaccine rollout programs.

Another compelling use case lies in humanitarian response systems, where the model could process crisis reports from multiple languages simultaneously, identifying urgent needs while respecting cultural expression norms. The low-memory footprint allows deployment on mobile devices used by field workers in connectivity-limited environments. The disentangled representations could further assist in distinguishing universal distress signals from culture-specific help-seeking behaviors.

Educational technology represents a third promising domain, where the quantized embeddings could power

personalized language learning tools that adapt to students' native language structures while maintaining cross-lingual semantic relationships. The efficient representation would enable offline use in resource-constrained educational settings, particularly valuable in developing regions with limited internet access.

6.3 Ethical Issues in Multilingual Embeddings

The development of multilingual technologies raises important ethical considerations that our framework only partially addresses. While the explicit separation of language-invariant and culture-specific features helps mitigate some forms of bias, the training process still relies on existing corpora that may underrepresent minority languages and dialects. Our experiments show performance gaps for low-resource languages that persist even after quantization-aware training, potentially exacerbating digital divides.

The cultural adaptation mechanism introduces additional concerns about representation fairness. The gating weights learned from majority language data may inadvertently suppress legitimate cultural expressions from marginalized communities when these differ significantly from dominant patterns. Furthermore, the quantization process itself could introduce new forms of bias by disproportionately compressing features important for certain cultural contexts.

Future work should investigate auditing frameworks that evaluate both the semantic and ethical impacts of quantization across different language communities. Developing fairness-aware quantization strategies that explicitly preserve features critical for marginalized groups could help address these concerns. The integration of community-based validation protocols, where native speakers assess the cultural appropriateness of quantized outputs, represents another important direction for ensuring responsible deployment.

7 Conclusion

The InfoBottleneck Quantized Multilingual Embeddings framework demonstrates that effective cross-cultural opinion analysis can be achieved without sacrificing computational efficiency. By reformulating multilingual representation learning through an information-theoretic lens, we maintain semantic alignment while reducing memory requirements by 4× compared to conventional approaches. The disentangled architecture preserves both language-invariant semantics and

culture-specific expressions, enabling more nuanced analysis of opinion shifts across linguistic boundaries.

Quantization-aware training guided by corpus linguistics principles ensures robust performance even at aggressive 4-bit precision, with only a 3.3% average accuracy drop on cross-lingual tasks. The dynamic gating mechanism adapts embeddings to event-specific cultural contexts without retraining, addressing a critical limitation in static multilingual models. Our experiments show consistent improvements over state-of-the-art baselines across semantic alignment, quantization robustness, and opinion analysis tasks.

The framework’s practical impact extends beyond academic benchmarks, offering tangible benefits for real-world applications. The custom CUDA kernels enable efficient deployment on resource-constrained devices, making cross-cultural opinion monitoring feasible in edge computing scenarios. The explicit separation of universal and culture-specific features provides interpretable insights into how opinions evolve differently across linguistic communities during global events.

Future extensions could explore adaptive quantization strategies that automatically adjust bit-widths based on language typology and task requirements. Integrating the framework with emerging techniques in few-shot learning may further enhance its applicability to low-resource languages. The principles demonstrated here suggest promising directions for developing efficient, culturally-aware NLP systems that respect linguistic diversity while maintaining practical deployability.

References:

- [1] C. Wang and M. Banko, “Practical transformer-based multilingual text classification,” in *Proceedings of* 2021.
- [2] U. Kulkarni, S. Meena, P. Joshua, K. Kruthika, *et al.*, “Performance improvements in quantization aware training and appreciation of low precision computation in deep learning,” *Unable to determine the complete publication venue*, 2020.
- [3] N. Tishby and N. Zaslavsky, “Deep learning and the information bottleneck principle,” in *2015 IEEE information theory workshop*, 2015.
- [4] B. Thompson, S. Roberts, *et al.*, “Quantifying semantic similarity across languages,” *Unable to Determine*, 2018.
- [5] I. Vulić and M. Moens, “Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings,” in *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 2015.
- [6] A. Conneau, K. Khandelwal, N. Goyal, *et al.*, “Unsupervised cross-lingual representation learning at scale,” arXiv preprint arXiv:1911.02116, 2019.
- [7] L. Wu, S. Wu, X. Zhang, D. Xiong, S. Chen, *et al.*, “Learning disentangled semantic representations for zero-shot cross-lingual transfer in multilingual machine reading comprehension,” arXiv preprint arXiv:2204.00996, 2022.
- [8] Z. Liu, Y. Wang, K. Han, W. Zhang, *et al.*, “Post-training quantization for vision transformer,” in *Advances in neural information processing systems*, 2021.
- [9] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, *et al.*, “Quantized neural networks: Training neural networks with low precision weights and activations,” *Journal of Machine Learning Research*, 2018.
- [10] A. Jiang, L. Du, and Y. Du, “Groupq: Group-wise quantization with multi-objective optimization for cnn accelerators,” in *IEEE/ACM international conference on computer-aided design*, 2024.
- [11] J. Li and D. Liu, “Information bottleneck theory on convolutional neural networks,” *Neural Processing Letters*, 2021.
- [12] A. See, M. Luong, and C. Manning, “Compression of neural machine translation models via pruning,” arXiv preprint arXiv:1606.09274, 2016.
- [13] C. Troussas and M. Virvou, “Information theoretic clustering for an intelligent multilingual tutoring system,” ... of Emerging Technologies in Learning, 2013.
- [14] K. Aung and N. Myo, “Sentiment analysis of students’ comment using lexicon based approach,” ... on computer; information science, 2017.
- [15] I. Moutidis and H. Williams, “Good and bad events: Combining network-based event detection with sentiment analysis,” *Social Network Analysis and Mining*, 2020.
- [16] Y. Choukroun, E. Kravchik, F. Yang, *et al.*, “Low-bit quantization of neural networks for efficient inference,” in *2019 IEEE/CVF conference on computer vision and*

- pattern recognition (CVPR)*, 2019.
- [17] S. Esser, J. McKinstry, D. Bablani, *et al.*, “Learned step size quantization,” arXiv preprint arXiv:1902.08153, 2019.
- [18] Z. Dong, Z. Yao, A. Gholami, *et al.*, “Hawq: Hessian aware quantization of neural networks with mixed-precision,” in *Proceedings of the IEEE international conference on computer vision*, 2019.
- [19] R. Müller, *Language universals in the brain: How linguistic are they*. Language universals, 2009.
- [20] R. Gray and D. Neuhoff, “Quantization,” *IEEE transactions on information theory*, 2002.
- [21] S. Lorenzen, C. Igel, and M. Nielsen, “Information bottleneck: Exact analysis of (quantized) neural networks,” arXiv preprint arXiv:2106.12912, 2021.
- [22] S. Malmasi and M. Dras, “Multilingual native language identification,” *Natural Language Engineering*, 2017.
- [23] S. Yu, F. Alesiani, X. Yu, R. Jenssen, *et al.*, “Measuring dependence with matrix-based entropy functional,” in *Proceedings of the association for the advancement of artificial intelligence*, 2021.
- [24] J. Piskorski and G. Jacquet, “TF-IDF character n-grams versus word embedding-based models for fine-grained event classification: A preliminary study,” in *Proceedings of the workshop on automated extraction of social and political events from news (AESPEN)*, 2020.
- [25] A. Conneau, G. Lample, R. Rinott, A. Williams, *et al.*, “XNLI: Evaluating cross-lingual sentence representations,” arXiv preprint arXiv:1809.05053, 2018.
- [26] J. Ham and E. Kim, “Semantic alignment with calibrated similarity for multilingual sentence embedding,” *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2021.
- [27] N. van der Heijden, H. Yannakoudakis, P. Mishra, *et al.*, “Multilingual and cross-lingual document classification: A meta-learning approach,” arXiv preprint arXiv:2101.11302, 2021.
- [28] E. Zotova, R. Aggeri, M. Nuñez, *et al.*, “Multilingual stance detection in tweets: The catalonia independence corpus,” in *Proceedings of the twelfth language resources and evaluation conference*, 2020.
- [29] T. Pires, E. Schlinger, and D. Garrette, “How multilingual is multilingual BERT?” arXiv preprint arXiv:1906.01502, 2019.
- [30] O. Zafrir, G. Boudoukh, P. Izsak, *et al.*, “Q8bert: Quantized 8bit bert,” in *2021 design, automation & test in europe conference & exhibition (date)*, 2019.
- [31] Z. Yao, R. Y. Aminabadi, *et al.*, “Zeroquant: Efficient and affordable post-training quantization for large-scale transformers,” in *Advances in neural information processing systems*, 2022.
- [32] J. Ying, W. Tang, Y. Zhao, Y. Cao, Y. Rong, *et al.*, “Disentangling language and culture for evaluating multilingual large language models,” arXiv preprint arXiv:2505.24635, 2025.
- [33] Z. Chi, L. Dong, F. Wei, N. Yang, S. Singhal, *et al.*, “InfoXLM: An information-theoretic framework for cross-lingual language model pre-training,” arXiv preprint arXiv:2007.07834, 2020.