

# GloCTM: Cross-Lingual Topic Modeling via a Global Context Space

Nguyen Tien Phat<sup>1\*</sup>, Ngo Vu Minh<sup>1\*</sup>, Linh Van Ngo<sup>1†</sup>, Nguyen Thi Ngoc Diep<sup>2</sup>,  
Thien Huu Nguyen<sup>3</sup>

<sup>1</sup> Hanoi University of Science and Technology, Vietnam

<sup>2</sup> VNU University of Engineering and Technology, Vietnam

<sup>3</sup> University of Oregon, US

{phat.nt235544, minh.nv230084}@sis.hust.edu.vn, linhnv@soict.edu.vn, ngocdiep@vnu.edu.vn, thienn@uoregon.edu

## Abstract

Cross-lingual topic modeling seeks to uncover coherent and semantically aligned topics across languages—a task central to multilingual understanding. Yet most existing models learn topics in disjoint, language-specific spaces and rely on alignment mechanisms (e.g., bilingual dictionaries) that often fail to capture deep cross-lingual semantics, resulting in loosely connected topic spaces. Moreover, these approaches often overlook the rich semantic signals embedded in multilingual pretrained representations, further limiting their ability to capture fine-grained alignment. We introduce **GloCTM** (**G**lobal **C**ontext **S**pace for **C**ross-Lingual **T**opic **M**odel), a novel framework that enforces cross-lingual topic alignment through a unified semantic space spanning the entire model pipeline. GloCTM constructs enriched input representations by expanding bag-of-words with cross-lingual lexical neighborhoods, and infers topic proportions using both local and global encoders, with their latent representations aligned through internal regularization. At the output level, the global topic-word distribution, defined over the combined vocabulary, structurally synchronizes topic meanings across languages. To further ground topics in deep semantic space, GloCTM incorporates a Centered Kernel Alignment (CKA) loss that aligns the latent topic space with multilingual contextual embeddings. Experiments across multiple benchmarks demonstrate that GloCTM significantly improves topic coherence and cross-lingual alignment, outperforming strong baselines.

**Code** — <https://github.com/tienphat140205/GloCTM>

## Introduction

Uncovering latent thematic structures from large-scale text corpora is a foundational task in natural language processing, with topic modeling (TM) serving as a powerful technique for distilling meaning from data (Blei, Ng, and Jordan 2003). The challenge intensifies in a globalized world, where understanding shared narratives across different languages is paramount. Cross-Lingual Topic Modeling (CLTM) addresses this challenge, as illustrated in Figure 1. The objective is to discover coherent topics that maintain

\*These authors contributed equally.

†Corresponding Author.

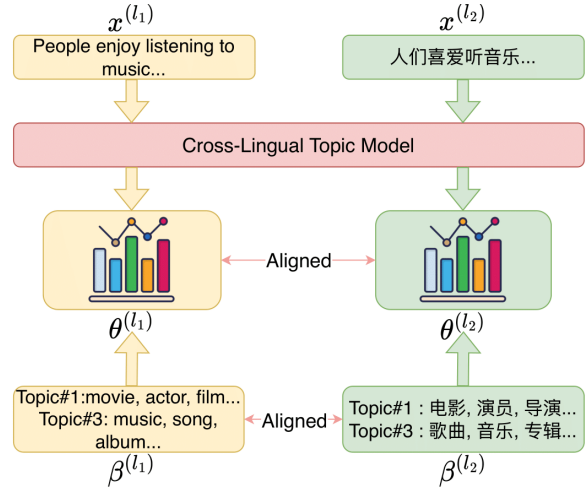


Figure 1: Conceptual overview of Cross-Lingual Topic Modeling (CLTM)

their semantic identity across linguistic boundaries (Mimno et al. 2009; Boyd-Graber and Blei 2009; Ni et al. 2009). For instance, given documents sharing the same semantic content across different languages, an ideal CLTM infers similar document-topic proportions ( $\theta$ ). This requires the underlying topics themselves, defined by their word distributions ( $\beta$ ), to be robustly aligned, such that a single topic index corresponds to the same concept in each language. This capability thereby enables large-scale cross-cultural analysis and global knowledge discovery.

Early efforts in cross-lingual topic modeling (CLTM) often relied on parallel corpora to learn aligned topics across languages (Mimno et al. 2009). However, such data is scarce and expensive to obtain, particularly for low-resource language pairs. As a result, most subsequent approaches turn to bilingual dictionaries as a more practical alternative (Jagaramudi and III 2010; Hu et al. 2014; Shi et al. 2016; Yuan, Durme, and Ying 2018; Wu et al. 2020, 2023). While effective to some extent, this strategy introduces two key challenges. First, dictionary coverage is often limited (Chang and Hwang 2021), leading to poor topic alignment (Jagaramudi and III 2010; Hao and Paul 2018). Second, accord-

<b>En Topic#1:</b>	game controller	golf	play	difficulty
<b>JP Topic#1:</b>	靴 かかと	つま先	履く	足
<b>Translation#1:</b>	shoes heel	toes	to wear	foot
<b>En Topic#2:</b>	collar honesty	gun	patience	operation
<b>JP Topic#2:</b>	首輪 チワワ	キロ	口コミ	家
<b>Translation#2:</b>	collar chihuahua	kg	review	house
<b>En Topic#3:</b>	shave shaver	trimmer	shampoo	gel
<b>JP Topic#3:</b>	歯 食い	知人	科	辛口
<b>Translation#3:</b>	tooth eat	acquaintance	department	dry taste

Table 1: Examples of misaligned topics generated by InfoCTM (Wu et al. 2023) across English (En) and Japanese (JP). The words grouped under each topic index show a clear semantic drift between the two languages.

ing to the analysis in (Wu et al. 2023), the direct alignment of topic-word distributions in models such as NMTM (Wu et al. 2020) and MCTA (Shi et al. 2016) can lead to a trivial solution. In this scenario, topic representations degenerate, causing the resulting topics to collapse and lose their interpretability. While recent models such as InfoCTM (Wu et al. 2023) have partially alleviated topic collapse and low dictionary coverage, their underlying architectural philosophy still inherits fundamental limitations from the broader CLTM landscape. A key concern is that these models continue to learn topic proportions ( $\theta$ ) and topic-word distributions ( $\beta$ ) in separate, disjoint spaces for each language, relying on an auxiliary loss function to indirectly bridge the gap between these otherwise disconnected language spaces. This architectural separation is the root cause of why topic alignment remains insufficiently robust. As illustrated in Figure 1, this frequently leads to a “semantic drift,” where a single topic index can correspond to a coherent theme in one language (e.g., “video games” in English) but a completely unrelated one in another (e.g., “footwear” in Japanese). This highlights the need for a framework that enforces alignment structurally, rather than handling it as an external or secondary component.

Moreover, many models still depend on static dictionaries with limited coverage and shallow lexical cues, while overlooking the rich semantic signals offered by multilingual pretrained language models. Recent works (Bianchi et al. 2021; Mueller and Dredze 2021) utilize multilingual BERT for topic modeling over multilingual corpora, but their goal is not to learn aligned cross-lingual topics, and thus they fall outside the scope of cross-lingual topic modeling. Another line of work, the clustering-based refinement method by (Chang et al. 2024), falls short, as it discovers latent groupings from multilingual embeddings but fails to model topic distributions, incorporate a generative structure, or enforce cross-lingual alignment. These trends highlight a critical gap: the absence of a unified framework that integrates multilingual embeddings with interpretable, aligned, and generative topic modeling.

To overcome these limitations, we introduce **GloCTM**, a dual-pathway neural topic model that brings topic alignment to the core of its design. Rather than relying on auxiliary losses to bridge disjoint language spaces, GloCTM con-

structs a shared semantic space from the outset. It begins by enriching each document with semantically related words from both languages (identified via multilingual pretrained embeddings) to form a dynamic global bag-of-words representation. This enriched input is passed through a global encoder, which infers a unified topic proportion vector ( $\theta_{\text{global}}$ ), later decoded via a global topic-word matrix ( $\beta_{\text{global}}$ ) that is constructed from the topic-word matrices of the individual language pathways to structurally enforce cross-lingual topic sharing. At the same time, language-specific local encoders capture monolingual subtleties, while a KL divergence loss encourages their outputs to stay aligned with the global topic view. This joint architecture allows GloCTM to balance language-specific expressiveness with semantic coherence across languages—offering a principled, end-to-end solution to cross-lingual topic modeling.

In light of the rapid advances in multilingual pretrained language models, incorporating their semantic knowledge into cross-lingual topic modeling has become not only natural but essential. These embeddings provide a rich, language-agnostic semantic space that can inform topic inference across languages. To harness this, we introduce a final calibration mechanism based on Centered Kernel Alignment (CKA) (Cortes, Mohri, and Rostamizadeh 2012; Kornblith et al. 2019). Specifically, we apply a CKA loss to align the model’s latent topic proportion vectors with the deep semantic representations produced by multilingual pretrained language models (PLMs). This regularization enhances the semantic structure of the topic proportions ( $\theta$ ), enabling them to reflect cross-lingual meaning beyond lexical overlap. As a result, **GloCTM** produces not only structurally aligned topics, but also semantically coherent ones as a natural consequence of improved latent representations. Our main contributions are summarized as follows:

- We introduce **GloCTM**, a novel dual-pathway framework for cross-lingual topic modeling that combines a global and local VAE architecture, while leveraging multilingual pretrained embeddings to guide topic learning across languages.
- We propose a **dynamic global bag-of-words construction** that enriches each document with semantically related words from both languages, enabling alignment-aware input representations grounded in lexical neighborhoods.
- We enforce robust alignment through a dual strategy: structurally, by decoding through a **unified topic-word matrix**, and representationally, through a **KL divergence loss** for internal consistency and a **Centered Kernel Alignment (CKA) loss** for semantic grounding.
- We empirically demonstrate that **GloCTM** consistently outperforms strong baselines on both topic diversity and cross-lingual alignment across multiple datasets and language pairs.

## Related Work

**Cross-Lingual Topic Modeling.** Topic models aim to uncover latent semantic structures in large-scale corpora, with

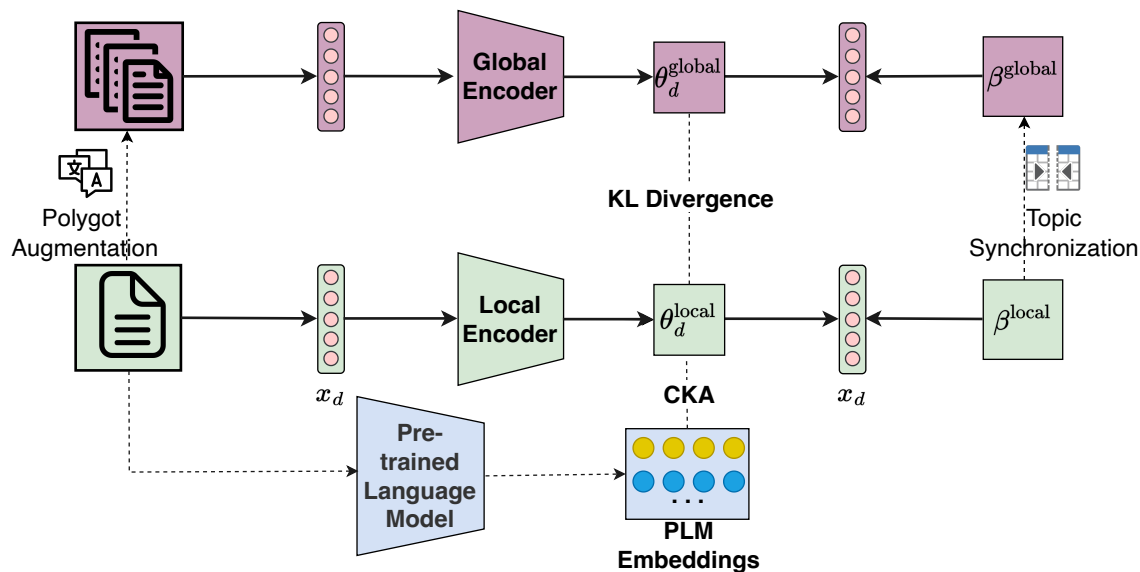


Figure 2: The dual-pathway architecture of GloCTM. The model integrates a Local Pathway for language-specific features with a Global Pathway for unified topic learning. The Global Pathway is uniquely fed by a semantically-enriched input created via Polyglot Augmentation. Alignment is enforced through three core mechanisms: Topic Synchronization ( $\beta$ ), KL Divergence ( $\theta_{local} \leftrightarrow \theta_{global}$ ), and CKA loss with PLM embeddings.

Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003) as a foundational example. Neural topic models (NTMs) extend this framework using variational autoencoders (e.g., NVDM, ProLDA (Srivastava and Sutton 2017)) and contextual embeddings (Dieng, Ruiz, and Blei 2020; Bianchi et al. 2021).

Cross-lingual topic modeling (CLTM) adapts this idea to multilingual data, seeking coherent and aligned topics across languages. Early CLTM relied on parallel corpora (Mimno et al. 2009), which are scarce. Dictionary-based methods became standard (Jagarlamudi and III 2010; Hu et al. 2014; Shi et al. 2016; Yuan, Durme, and Ying 2018; Yang, Boyd-Graber, and Resnik 2017, 2019; Wu et al. 2020), but suffer from limited coverage, weak semantics, and topic collapse. InfoCTM (Wu et al. 2023) improves topic diversity via contrastive learning and introduces a novel dictionary expansion technique, Cross-lingual Vocabulary Linking (CVL). Recently, the XTRA framework (Nguyen et al. 2025) combines bag-of-words modelling with multilingual embeddings and dual contrastive alignment of document-topic and topic-word distributions to yield interpretable and well-aligned cross-lingual topics. Other work uses pretrained transformers for zero-shot modeling (Bianchi et al. 2021; Mueller and Dredze 2021), though lacking mechanisms for interpretable topic-word distributions. Refining Dimensions (Chang et al. 2024), based on BERTopic (Grootendorst 2022), improves alignment by removing noisy embedding dimensions but does not infer topic or document distributions. These limitations highlight a key gap: current CLTM approaches often trade off interpretability or rely on shallow alignment, motivating the need for a principled, generative model that ensures both semantic depth and cross-lingual consistency.

**Centered Kernel Alignment (CKA).** CKA, introduced

in kernel learning (Cortes, Mohri, and Rostamizadeh 2012), extends the Hilbert-Schmidt Independence Criterion by computing normalized similarity between centered Gram matrices. Popularized in deep learning by Kornblith et al. (Kornblith et al. 2019), it captures representation similarity across layers and architectures. CKA has since been used in transfer learning (Bonheme and Grzes 2023; Matsoukas et al. 2022), architecture analysis (Raghu et al. 2021), knowledge distillation (Saha, Bialkowski, and Khalifa 2022; Zhou et al. 2024), and model pruning (Pons et al. 2024). Despite its potential, CKA remains underexplored in topic modeling, where it could align document-topic and embedding-based representations across languages.

## Methodology

We propose **GloCTM**, a dual-pathway Variational Autoencoder (VAE) framework comprising parallel local and global branches. The local branch consists of two language-specific VAEs—one for each language—while the global branch is shared and receives inputs enriched via our **Polyglot Augmentation** method. As illustrated in Figure 2, the model’s robustness arises from three core mechanisms: **Topic Synchronization** through a unified decoder, a **KL Divergence** loss for internal consistency, and a **Centered Kernel Alignment (CKA)** loss for semantic grounding in pre-trained embeddings. The following sections detail these components.

## Notations

We consider a multilingual corpus  $\mathcal{C}$  with documents in two languages,  $L_1$  and  $L_2$ , and aim to discover  $K$  shared topics. Each document  $d$  in language  $L_l$  is represented by: (i) a BoW vector  $\mathbf{x}_d^{(l)} \in \mathbb{R}^{|V^{(l)}|}$ , (ii) a cross-lingually enriched

vector  $\mathbf{g}_d^{(l)} \in \mathbb{R}^{|V^{(1)}|+|V^{(2)}|}$ , and (iii) an  $M$ -dimensional pre-trained embedding  $\mathbf{e}_d^{(l)}$ . Parallel local and global encoders map these inputs to latent topic proportions,  $\theta_d^{(l,\text{local})}$  and  $\theta_d^{(l,\text{global})}$ , which are decoded via language-specific matrices  $\beta^{(l)}$  and a unified matrix  $\beta^{(\text{global})}$  obtained by concatenating them.

### Polyglot Augmentation for a Pre-Aligned Global Context

A key limitation of traditional cross-lingual topic models is their use of language-specific BoW vectors that reside in disjoint feature spaces. Documents from different languages that convey the same semantic content often have non-overlapping vocabularies, forcing models to exert a corrective pressure through integrated loss functions, in a fragile attempt to harmonize their separate topic-word distributions.

To address this, we introduce *Polyglot Augmentation*, a method that constructs a shared **Global Context Space** directly at the input level by enriching each document’s Bag-of-Words (BoW) vector with both intra- and cross-lingual neighbors. Given a document-level BoW vector  $\mathbf{x}_d^{(l)} \in \mathbb{R}^{|V^{(l)}|}$  and its active word indices  $W_d^{(l)} = \{w_i : x_{d,i}^{(l)} > 0\}$ , we augment the representation as follows. First, we retrieve for each active word  $w \in W_d^{(l)}$  its top- $k$  most similar intra-lingual neighbors  $N_I(w)$  using cosine similarity over word embeddings in language  $l$ , and cross-lingual neighbors  $N_C(w)$  using similarity between embeddings across languages. The enriched BoW vectors are then defined as:

$$\tilde{x}_{d,j}^{(l)} = \sum_{w \in W_d^{(l)}} \left( [w_j^{(l)} = w] + [w_j^{(l)} \in N_I(w)] \right)$$

$$\tilde{x}_{d,j}^{(-l)} = \sum_{w \in W_d^{(l)}} [w_j^{(-l)} \in N_C(w)]$$

where  $[\cdot]$  denotes the Iverson bracket, which evaluates to 1 if the condition holds, and 0 otherwise. Finally, the global input vector is the concatenation of the intra- and cross-lingual components in a fixed layout:

$$\mathbf{g}_d^{(l)} = \begin{cases} [\tilde{\mathbf{x}}_d^{(l)} \mid \tilde{\mathbf{x}}_d^{(-l)}], & \text{if } l = \text{lang}_1 \\ [\tilde{\mathbf{x}}_d^{(-l)} \mid \tilde{\mathbf{x}}_d^{(l)}], & \text{if } l = \text{lang}_2 \end{cases}$$

This yields a cross-lingually comparable representation in the joint vocabulary space  $\mathbb{R}^{|V^{(1)}|+|V^{(2)}|}$ , ensuring that semantically related documents share overlapping features even across languages. This process is illustrated in Figure 3.

This construction yields document vectors that are cross-lingually comparable. For example, documents in  $\text{lang}_1$  and  $\text{lang}_2$  about the same topic (e.g., “football”) will include overlapping features such as soccer, goal, stadium, and player. Consequently, the global encoder no longer needs to infer alignment—it directly observes semantic proximity in a unified input space. This transformation turns alignment from a challenge of bridging separate language spaces into an intrinsic property of the model’s input.

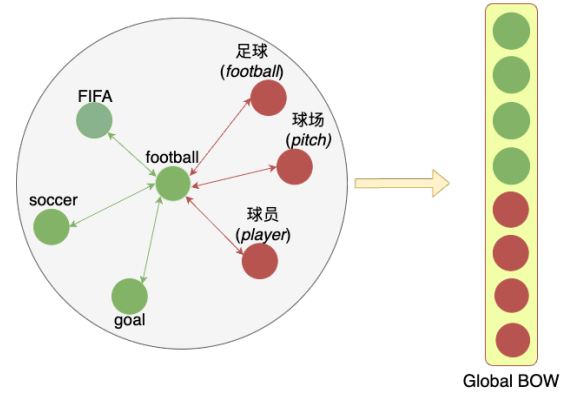


Figure 3: The Polyglot Augmentation mechanism. An original word (“football”) is enriched with its nearest intra-lingual (green) and cross-lingual (red) neighbors. This creates a dense Global BOW vector, preemptively injecting cross-lingual information before the encoding step.

### Internal Consistency via KL Divergence

Our dual-pathway architecture, with its language-specific **local encoders** and shared **global encoder**, risks learning divergent latent spaces that could compromise cross-lingual consistency. To enforce internal alignment, we introduce a regularization term,  $\mathcal{L}_{KL}$ , that minimizes the Kullback-Leibler (KL) divergence between the local and global posterior distributions for each document:

$$\mathcal{L}_{KL} = KL(q(z_d^{(l,\text{local})} | x_d^{(l)}) \parallel q(z_d^{(l,\text{global})} | g_d^{(l)}))$$

This loss serves a dual purpose. Primarily, it pulls the language-specific representations ( $\theta_{\text{local}}$ ) towards the shared semantic space defined by the global encoder. This ensures the encoders learn meaningful and high-quality topic proportion representations ( $\theta$ ), providing a crucial counterbalance to the strong structural constraints imposed by our unified decoder design.

### Topic Synchronization via a Unified Decoder

A core innovation of GLOCTM lies in its direct structural enforcement of topic alignment. Instead of applying alignment losses to indirectly connect separate language pathways, the model builds alignment directly into the decoder’s parameterization

Specifically, while language-specific topic-word matrices  $\beta^{(l)} \in \mathbb{R}^{K \times |V^{(l)}|}$  are used for local views, the global decoder constructs a unified topic-word matrix via horizontal concatenation:

$$\beta^{(\text{global})} = [\beta^{(1)} \mid \beta^{(2)}]$$

This design ensures that each topic  $k$  is represented by a single, contiguous vector spanning the joint vocabulary, inherently coupling the two languages. An example of this mechanism is illustrated in Figure 4.

The enriched input vector  $g_d$  contains words from both vocabularies, but all of them describe the same document and topic. As a result, the global decoder must reconstruct a




Figure 4: Formation of a unified ‘Food’ topic (Topic #2) in GloCTM. The global topic vector is constructed by concatenating the local topic vectors from English ( $\beta_2^{(1)}$ ) and Chinese ( $\beta_2^{(2)}$ ), enforcing semantic alignment by design.

multilingual signal that consistently represents a single underlying theme. This is achieved using a shared topic vector  $\theta_{\text{global}}$  and the unified matrix  $\beta^{(\text{global})}$ , where each row spans both vocabularies. To minimize reconstruction loss, the decoder is compelled to assign high weights to topically relevant words in *both* languages within the same topic row. If the two halves of a row represent different semantics—e.g., “Food” in one language and “Sports” in the other—reconstruction will fail, enforcing strong alignment across languages as a structural necessity.

### Semantic Knowledge Distillation via Centered Kernel Alignment (CKA)

The rise of multilingual pre-trained language models (MPLMs) provides an opportunity to infuse topic models with deep semantic knowledge. Although our architecture enforces structural alignment, the topic proportion vectors may lack semantic grounding. To enhance their expressiveness, we distill information from MPLMs. However, the mismatch between the  $K$ -dimensional topic space and the  $M$ -dimensional embedding space makes direct similarity comparison ineffective.

We therefore employ **Centered Kernel Alignment (CKA)** (Cortes, Mohri, and Rostamizadeh 2012; Kornblith et al. 2019), a robust metric ideal for comparing such disparate representations because it aligns the underlying geometry of the spaces rather than the raw vectors. The distillation objective minimizes the CKA loss between the inferred local topic proportion matrix ( $\Theta$ ) and the MPLM embedding matrix ( $E$ ):

$$\mathcal{L}_{\text{CKA}} = 1 - \text{CKA}(\Theta, E)$$

where CKA is the normalized Hilbert-Schmidt Independence Criterion (HSIC) between the Gram matrices of the two representations ( $\mathbf{K} = \Theta\Theta^T$ ,  $\mathbf{L} = EE^T$ ):

$$\text{CKA}(\Theta, E) = \frac{\text{HSIC}(\Theta, E)}{\sqrt{\text{HSIC}(\Theta, \Theta)\text{HSIC}(E, E)}}$$

This process infuses our topics with high-level semantic knowledge, ensuring they are interpretable and meaningfully grounded across languages.

### Training Objective

**VAE Objective per Pathway.** Each pathway in our dual-pathway framework is an instance of a VAE-based topic model. Following the standard Evidence Lower Bound (ELBO) objective (Kingma and Welling 2013), the loss for a single pathway processing an input  $\mathbf{x}$  is formulated as:

$$\mathcal{L}_{\text{VAE}}(\mathbf{x}) = -\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] + \text{KL}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))$$

The first term measures the reconstruction error (negative log-likelihood), while the second term is the KL divergence that regularizes the approximate posterior  $q(\mathbf{z}|\mathbf{x})$  to be close to the prior  $p(\mathbf{z})$ .

**Overall Objective for GloCTM.** The final training objective integrates the VAE losses from the global and local pathways, along with two alignment regularizers,  $\mathcal{L}_{\text{KL}}$  and  $\mathcal{L}_{\text{CKA}}$ . Letting  $\Phi$  denote all model parameters, the full model is trained end-to-end by minimizing.

$$\min_{\Phi} \mathcal{L} = \mathcal{L}_{\text{VAE}}^{(\text{global})} + \sum_{l \in \{1,2\}} \mathcal{L}_{\text{VAE}}^{(l, \text{local})} + \lambda_1 \mathcal{L}_{\text{KL}} + \lambda_2 \mathcal{L}_{\text{CKA}}$$

## Experiment

### Dataset

We evaluated our models on three benchmark datasets: EC News (Wu et al. 2020), containing English–Chinese news articles across six categories; Amazon Review (Yuan, Durme, and Ying 2018), an English–Chinese review dataset reformulated as a binary classification task (five-star = 1, others = 0); and Rakuten Amazon (Yuan, Durme, and Ying 2018), comprising Japanese Rakuten and English Amazon reviews, also cast as a binary rating-based classification task.

### Evaluation Metrics

We evaluated topic quality and utility using a multi-faceted approach. Intrinsic quality was measured by cross-lingual coherence (CNPMI) (Hao, Boyd-Graber, and Paul 2018; Chang et al. 2009) and topic diversity (TU) (Nan et al. 2019), with both metrics computed using the top 15 words of each topic, along with their resulting Topic Quality (TQ) (Dieng, Ruiz, and Blei 2020). Practical utility was assessed via SVM classification in intra-lingual (-I) and cross-lingual (-C) settings using the topic distributions as features, following prior work (Yuan, Durme, and Ying 2018; Wu et al. 2023). Finally, inspired by Stambach et al. (2023), we used Large Language Model (LLM) ratings for a human-centered evaluation of coherence and alignment.

### Baseline Models

We compare GloCTM with representative baselines: MCTA (Shi et al. 2016) (probabilistic cross-lingual topic modeling), MTAnchor (Yuan, Durme, and Ying 2018) (anchor-word alignment), NMTM (Wu et al. 2020) (neural shared-vocabulary alignment), and InfoCTM (Wu et al. 2023) (mutual-information-based alignment). We also include u-SVD and SVD-LR (Chang et al. 2024), clustering methods on multilingual embeddings, and XTRA (Nguyen et al. 2025), a cross-lingual topic model improving coherence and alignment.

Model	EC News			Amazon Review			Rakuten Amazon		
	CNPMI	TU	TQ	CNPMI	TU	TQ	CNPMI	TU	TQ
MCTA	0.025	0.489	0.012	0.028	0.319	0.009	0.021	0.272	0.006
MTAnchor	-0.013	0.192	0.000	0.028	0.323	0.009	-0.001	0.214	0.000
NMTM	0.028±0.004	0.840±0.016	0.023±0.004	0.043±0.003	0.673±0.028	0.029±0.003	0.010±0.001	0.706±0.022	0.007±0.001
InfoCTM	0.045±0.002	0.926±0.005	0.041±0.002	0.036±0.004	0.940±0.010	0.034±0.004	<u>0.033±0.001</u>	0.861±0.008	0.028±0.001
u-SVD	0.082±0.003	0.830±0.006	0.068±0.003	0.055±0.002	0.634±0.008	0.035±0.001	0.027±0.001	0.571±0.010	0.015±0.000
SVD-LR	<b>0.083±0.007</b>	0.820±0.007	<u>0.068±0.006</u>	0.053±0.003	0.627±0.013	0.034±0.002	0.026±0.002	0.558±0.010	0.015±0.001
XTRA	0.072±0.002	<u>0.982±0.002</u>	<b>0.070±0.002</b>	0.052±0.003	<b>0.976±0.003</b>	<u>0.050±0.003</u>	0.029±0.002	<b>0.971±0.007</b>	0.027±0.002
<b>GloCTM</b>	0.071±0.004	<b>0.985±0.009</b>	<b>0.070±0.004</b>	<b>0.058±0.002</b>	<u>0.958±0.004</u>	<b>0.056±0.002</b>	<b>0.040±0.004</b>	<u>0.925±0.006</u>	<b>0.037±0.002</b>

Table 2: Topic quality scores (CNPMI for coherence, TU for diversity, and  $TQ = \max(\text{CNPMI}, 0) \times TU$  across three datasets. Bold and underline indicate best and second-best scores, respectively.

Model	Topic Quality		Classification			
	CNPMI	TU	EN-I	ZH-I	EN-C	ZH-C
NMTM	0.045	0.643	0.788	0.721	0.592	0.575
InfoCTM	0.036	0.934	0.788	0.725	0.672	0.601
w/o $\mathcal{L}_{KL}$	0.058	0.949	0.803	0.722	0.708	0.640
w/ $\mathcal{L}_{sim}$	<b>0.059</b>	0.942	0.808	0.725	0.758	0.641
w/o $\mathcal{L}_{CKA}$	0.054	<b>0.961</b>	0.807	0.721	0.718	0.638
<b>GloCTM</b>	0.058	0.958	<b>0.818</b>	<b>0.728</b>	<b>0.763</b>	<b>0.642</b>

Table 3: Ablation study of GloCTM on Amazon Review. Scores are means of five seeds; column bests are in bold.

### Overall Topic Quality

As shown in Table 2, GloCTM is overall dominant compared to the baselines in terms of topic quality—assessed by coherence (CNPMI), diversity (Topic Uniqueness, TU), and the combined Topic Quality (TQ) score—consistently achieving the highest TU and TQ scores across all datasets. For instance, on the Rakuten Amazon dataset, its CNPMI (0.040) and TQ (0.037) are markedly higher than the best-performing baselines, InfoCTM (CNPMI = 0.033) and XTRA (TQ = 0.027). Moreover, its qualities (TQ) remain top-tier, achieving the highest scores across all three datasets—EC News (0.070, outperforming InfoCTM at 0.041), Amazon Review (0.056, surpassing XTRA at 0.050), and Rakuten Amazon (0.037, exceeding InfoCTM at 0.028). This strong performance across metrics reinforces GloCTM’s lead. Although TU drops slightly, this mainly comes from minor overlaps in key meaningful words—an expected effect of GloCTM’s tighter semantic vocabularies.

### Evaluation of Topic Distributions via Classification

To evaluate the quality and transferability of document-topic distributions, we use each model’s topic vectors  $\theta$  as features for SVM-based text classification, considering both intra-lingual (-I) setups (train/test in the same language) and cross-lingual (-C) setups (train on one language, test on another), as shown in Figure 5. GloCTM remains competitive





Figure 5: Classification performance on the three datasets. Except for MCTA and MTAnchor (results reported in (Wu et al. 2023)), all scores are averaged over 5 runs.


with classification-focused baselines such as XTRA, showing a slight but consistent advantage over the remaining baselines in intra-lingual settings (EN-I, ZH-I, JA-I), and a clearly superior performance in cross-lingual settings (EN-C, ZH-C, JA-C). Across all three datasets—EC News, Amazon Review, and Rakuten Amazon—GloCTM surpasses every baseline by a substantial margin in cross-lingual transfer.

### LLM-Assisted Evaluation of Topic Coherence and Alignment


Following recent work on LLM-based topic evaluation (Stambach et al. 2023), we employ large language models to assess cross-lingual semantic similarity between aligned topic pairs, using a 1–3 scoring scale averaged over four independent runs; the final reported score for each topic pair is then rounded to the nearest integer to ensure consistency across evaluations. As shown in Figure 6, on the two English–Chinese datasets, GloCTM achieves performance comparable to strong baselines such as SVD-LR and u-SVD while exhibiting greater stability on lower-quality topics, indicating more reliable semantic alignment. The distinction becomes more pronounced on the English–Japanese dataset, where the larger linguistic gap raises alignment difficulty: GloCTM produces a substantially higher number of strong, well-aligned topics and avoids the sharp performance fluctuations observed in competing models. Overall, these LLM-based assessments underscore both the robustness and the cross-lingual reliability of GloCTM’s topic representations.



(a) Cross-lingual semantic similarity on EC News dataset.



(b) Cross-lingual semantic similarity on Amazon Review dataset.



(c) Cross-lingual semantic similarity on Rakuten Amazon dataset.

Figure 6: LLM-based cross-lingual semantic similarity on Amazon Review, EC News, and Rakuten-Amazon; darker shades indicate higher scores (1–3). Values are means of four runs, rounded.

## Ablation Study

To validate our design, we perform an ablation study on the Amazon Review dataset, targeting two key losses: the alignment loss  $\mathcal{L}_{KL}$  and the structural loss  $\mathcal{L}_{CKA}$ . We evaluate three variants: (1) removing  $\mathcal{L}_{CKA}$ , (2) removing  $\mathcal{L}_{KL}$ , and (3) replacing  $\mathcal{L}_{KL}$  with a cosine similarity loss  $L_{sim}$ . As shown in Table 3, removing  $\mathcal{L}_{CKA}$  degrades classification performance (e.g., EN-C drops to 0.718), confirming its role in latent structure refinement. More critically, removing  $\mathcal{L}_{KL}$  leads to a substantial drop in classification accuracy (EN-C falls to 0.708), highlighting its vital importance in aligning topic distributions. Replacing  $\mathcal{L}_{KL}$  with  $\mathcal{L}_{sim}$  is also observed to be a less effective strategy, as it slightly lowers classification accuracy. These results demonstrate that GloCTM’s effectiveness stems from the synergy between a strong alignment signal ( $\mathcal{L}_{KL}$ ) and a structure-enforcing signal ( $\mathcal{L}_{CKA}$ ), leading to the most robust overall performance.

## Interpreting the Discovered Topics

As shown in Figure 4, GloCTM generates cleaner, better-aligned topics than InfoCTM, SVD-LR, and u-SVD. In the Media and Content topic, for instance, InfoCTM includes unrelated terms like “sky” and “dream”, SVD-LR suffers from noise and misalignment, and u-SVD mixes in sentiment. In contrast, GloCTM produces focused themes,

Topic: Media and Content					
InfoCTM Topic 19					
EN Topic:	season	sky	release	episodes	theatrical
ZH Topic:	星星	春天	梦想	续集	星空
Translations:	stars	spring	dream	sequel	starry sky
SVD-LR Topic 24					
EN Topic:	premiere	trailer	one	debut	first
ZH Topic:	第一	比赛	一一	之后	一月
Translations:	first	competition	one by one	afterwards	January
u-SVD Topic 21					
EN Topic:	movie	book	boring	read	characters
ZH Topic:	烂烂	失望	第三	内容	感觉
Translations:	lousy	disappointed	third	content	feeling
GloCTM Topic 45					
EN Topic:	cinema	movies	film	movie	films
ZH Topic:	电影院	影院	电影	影视	本片
Translations:	cinema	movie theater	movie	film/tv	this film
Topic: Product Quality and Materials					
InfoCTM Topic 39					
EN Topic:	seal	shell	housing	glue	fake
ZH Topic:	密封	暗	露	包装盒	盖子
Translations:	sealed	dark	reveal	packaging box	lid
SVD-LR Topic 53					
EN Topic:	review	display	guns	flip	new
ZH Topic:	采用	相机	英寸	像素	零零
Translations:	use/adopt	camera	inch	pixel	zero zero
u-SVD Topic 7					
EN Topic:	product	plastic	one	would	bag
ZH Topic:	卓越	产品	收到	包装	东西
Translations:	excellent	product	received	packaging	thing
GloCTM Topic 32					
EN Topic:	timely	condition	arrival	arrived	arrive
ZH Topic:	很快	塑料袋	盒子	完好	纸箱
Translations:	very fast	plastic bag	box	intact	carton

Table 4: Cross-lingual topic comparison for InfoCTM, SVD-LR, u-SVD, and GloCTM; red words indicate noise or misalignment.

with English keywords (e.g., “cinema,” “film”) aligned with their Chinese counterparts. Similarly, in the Product Quality topic, GloCTM offers coherent terms (e.g., “timely”, “arrival”), while baselines—especially SVD-LR—introduce off-topic words like “guns” and show topic repetition.

## Conclusion

We propose GloCTM, a dual-pathway framework enforcing intrinsic cross-lingual alignment through Polyglot Augmentation, a shared decoder, and alignment losses (KL, CKA). Experiments show that GloCTM outperforms baselines in topic coherence, diversity, and cross-lingual classification, demonstrating the effectiveness of its alignment mechanism.

## Acknowledgements

Linh Ngo Van is funded by Vietnam National Foundation for Science and Technology Development (NAFOSTED) un-

der grant number 102.05-2025.16. Thien Huu Nguyen has been supported by the NSF grant # 2239570. He is also supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the HIATUS Program contract 2022-22072200003.

## References

- Bianchi, F.; Terragni, S.; Hovy, D.; Nozza, D.; and Fersini, E. 2021. Cross-lingual Contextualized Topic Models with Zero-shot Learning. In Merlo, P.; Tiedemann, J.; and Tsarfaty, R., eds., *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, 1676–1683. Association for Computational Linguistics.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 993–1022.
- Bonheme, L.; and Grzes, M. 2023. How good are variational autoencoders at transfer learning? *arXiv preprint arXiv:2304.10767*.
- Boyd-Graber, J. L.; and Blei, D. M. 2009. Multilingual Topic Models for Unaligned Text. In Bilmes, J. A.; and Ng, A. Y., eds., *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*, 75–82. AUAI Press.
- Chang, C.; Huang, T.; Tsai, Y.; Chang, C.; and Hwang, S. 2024. Refining Dimensions for Improving Clustering-based Cross-lingual Topic Models. *CoRR*, abs/2412.12433.
- Chang, C.; and Hwang, S. 2021. A word embedding-based approach to cross-lingual topic modeling. *Knowl. Inf. Syst.*, 63(6): 1529–1555.
- Chang, J. D.; Boyd-Graber, J. L.; Gerrish, S.; Wang, C.; and Blei, D. M. 2009. Reading Tea Leaves: How Humans Interpret Topic Models. In Bengio, Y.; Schuurmans, D.; Lafferty, J. D.; Williams, C. K. I.; and Culotta, A., eds., *Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems 2009. Proceedings of a meeting held 7-10 December 2009, Vancouver, British Columbia, Canada*, 288–296. Curran Associates, Inc.
- Cortes, C.; Mohri, M.; and Rostamizadeh, A. 2012. Algorithms for Learning Kernels Based on Centered Alignment. *J. Mach. Learn. Res.*, 13: 795–828.
- Dieng, A. B.; Ruiz, F. J.; and Blei, D. M. 2020. Topic Modeling in Embedding Spaces. *Transactions of the Association for Computational Linguistics*, 439–453.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv:2203.05794*.
- Hao, S.; Boyd-Graber, J. L.; and Paul, M. J. 2018. Lessons from the Bible on Modern Topics: Low-Resource Multilingual Topic Model Evaluation. In Walker, M. A.; Ji, H.; and Stent, A., eds., *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, 1090–1100. Association for Computational Linguistics.
- Hao, S.; and Paul, M. J. 2018. Learning Multilingual Topics from Incomparable Corpora. In Bender, E. M.; Derczynski, L.; and Isabelle, P., eds., *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, 2595–2609. Association for Computational Linguistics.
- Hu, Y.; Zhai, K.; Eidelman, V.; and Boyd-Graber, J. L. 2014. Polylingual Tree-Based Topic Models for Translation Domain Adaptation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, 1166–1176. The Association for Computer Linguistics.
- Jagarlamudi, J.; and III, H. D. 2010. Extracting Multilingual Topics from Unaligned Comparable Corpora. In Gurrin, C.; He, Y.; Kazai, G.; Kruschwitz, U.; Little, S.; Roelleke, T.; R ger, S. M.; and van Rijsbergen, K., eds., *Advances in Information Retrieval, 32nd European Conference on IR Research, ECIR 2010, Milton Keynes, UK, March 28-31, 2010. Proceedings*, volume 5993 of *Lecture Notes in Computer Science*, 444–456. Springer.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations*.
- Kornblith, S.; Norouzi, M.; Lee, H.; and Hinton, G. E. 2019. Similarity of Neural Network Representations Revisited. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 3519–3529. PMLR.
- Matsoukas, C.; et al. 2022. What Makes Transfer Learning Work for Medical Images: Feature Reuse & Other Factors. In *CVPR*.
- Mimno, D. M.; Wallach, H. M.; Naradowsky, J.; Smith, D. A.; and McCallum, A. 2009. Polylingual Topic Models. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6-7 August 2009, Singapore, A meeting of SIGDAT, a Special Interest Group of the ACL*, 880–889. ACL.
- Mueller, A.; and Dredze, M. 2021. Fine-tuning Encoders for Improved Monolingual and Zero-shot Polylingual Neural Topic Modeling. In Toutanova, K.; Rumshisky, A.; Zettlemoyer, L.; Hakkani-T r, D.; Beltagy, I.; Bethard, S.; Cotterell, R.; Chakraborty, T.; and Zhou, Y., eds., *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, 3054–3068. Association for Computational Linguistics.
- Nan, F.; Ding, R.; Nallapati, R.; and Xiang, B. 2019. Topic Modeling with Wasserstein Autoencoders. In Korhonen, A.; Traum, D. R.; and M rquez, L., eds., *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019*,

- Volume 1: Long Papers*, 6345–6381. Association for Computational Linguistics.
- Nguyen, T. P.; Minh, N. V.; Nguyen, T.; Van, L. N.; Nguyen, D. A.; Sang, D. V.; and Le, T. 2025. XTRA: Cross-Lingual Topic Modeling with Topic and Representation Alignments. In Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; and Peng, V., eds., *Findings of the Association for Computational Linguistics: EMNLP 2025*, 5561–5575. Suzhou, China: Association for Computational Linguistics. ISBN 979-8-89176-335-7.
- Ni, X.; Sun, J.; Hu, J.; and Chen, Z. 2009. Mining multilingual topics from wikipedia. In Quemada, J.; León, G.; Maarek, Y. S.; and Nejd, W., eds., *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, 1155–1156. ACM.
- Pons, I.; Yamamoto, B.; Costa, A. H. R.; and Jordão, A. 2024. Effective Layer Pruning Through Similarity Metric Perspective. In Antonacopoulos, A.; Chaudhuri, S.; Chellappa, R.; Liu, C.; Bhattacharya, S.; and Pal, U., eds., *Pattern Recognition - 27th International Conference, ICPR 2024, Kolkata, India, December 1-5, 2024, Proceedings, Part V*, volume 15305 of *Lecture Notes in Computer Science*, 423–438. Springer.
- Raghu, M.; Unterthiner, T.; Kornblith, S.; Zhang, C.; and Dosovitskiy, A. 2021. Do Vision Transformers See Like Convolutional Neural Networks? In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 12116–12128.
- Saha, A.; Bialkowski, A.; and Khalifa, S. 2022. Distilling Representational Similarity using Centered Kernel Alignment (CKA). In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*, 535. BMVA Press.
- Shi, B.; Lam, W.; Bing, L.; and Xu, Y. 2016. Detecting Common Discussion Topics Across Culture From News Reader Comments. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics.
- Srivastava, A.; and Sutton, C. 2017. Autoencoding variational inference for topic models. In *International Conference on Learning Representations*.
- Stammbach, D.; Zouhar, V.; Hoyle, A. M.; Sachan, M.; and Ash, E. 2023. Revisiting Automated Topic Model Evaluation with Large Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, 9348–9357. Association for Computational Linguistics.
- Wu, X.; Dong, X.; Nguyen, T.; Liu, C.; Pan, L.; and Luu, A. T. 2023. InfoCTM: A Mutual Information Maximization Perspective of Cross-Lingual Topic Modeling. In Williams, B.; Chen, Y.; and Neville, J., eds., *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, 13763–13771. AAAI Press.
- Wu, X.; Li, C.; Zhu, Y.; and Miao, Y. 2020. Learning Multilingual Topics with Neural Variational Inference. In Zhu, X.; Zhang, M.; Hong, Y.; and He, R., eds., *Natural Language Processing and Chinese Computing - 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14-18, 2020, Proceedings, Part I*, volume 12430 of *Lecture Notes in Computer Science*, 840–851. Springer.
- Yang, W.; Boyd-Graber, J. L.; and Resnik, P. 2017. Adapting Topic Models using Lexical Associations with Tree Priors. In Palmer, M.; Hwa, R.; and Riedel, S., eds., *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, 1901–1906. Association for Computational Linguistics.
- Yang, W.; Boyd-Graber, J. L.; and Resnik, P. 2019. A Multilingual Topic Model for Learning Weighted Topic Links Across Corpora with Low Comparability. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 1243–1248. Association for Computational Linguistics.
- Yuan, M.; Durme, B. V.; and Ying, J. L. 2018. Multilingual Anchoring: Interactive Topic Modeling and Alignment Across Languages. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 8667–8677.
- Zhou, Z.; Shen, Y.; Shao, S.; Gong, L.; and Lin, S. 2024. Rethinking Centered Kernel Alignment in Knowledge Distillation. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, 5680–5688. ijcai.org.