

From FusHa to Folk: Exploring Cross-Lingual Transfer in Arabic Language Models

Abdulmuizz Khalak Abderrahmane Issam Gerasimos Spanakis

Department of Advanced Computing Sciences

Maastricht University

{a.khalak@student., abderrahmane.issam@, jerry.spanakis@}maastrichtuniversity.nl

Abstract

Arabic Language Models (LMs) are pretrained predominately on Modern Standard Arabic (MSA) and are expected to transfer to its dialects. While MSA as the standard written variety is commonly used in formal settings, people speak and write online in various dialects that are spread across the Arab region. This poses limitations for Arabic LMs, since its dialects vary in their similarity to MSA. In this work we study cross-lingual transfer of Arabic models using probing on 3 Natural Language Processing (NLP) Tasks, and representational similarity. Our results indicate that transfer is possible but disproportionate across dialects, which we find to be partially explained by their geographic proximity. Furthermore, we find evidence for negative interference in models trained to support all Arabic dialects. This questions their degree of similarity, and raises concerns for cross-lingual transfer in Arabic models.¹

1 Introduction

Arabic is a major world language with over 400 million speakers and a central role in cultural, historical, and religious life across the Middle East and North Africa (MENA) (Owens, 2006; Versteegh, 2014). Structurally, it is a *Semitic* language with a long written tradition, but for NLP the most salient property is its pervasive *diglossia*: a standardized written variety, Modern Standard Arabic (MSA), coexists with a rich spectrum of non-standardized dialects that vary across the MENA region (Ferguson, 1959). MSA (FusHa) dominates education, news, and formal writing, whereas dialects are the default medium of everyday communication and online interaction.

Although Modern Standard Arabic (MSA) is far less frequent in daily speech than Dialectal Arabic (DA), it remains the dominant variety in the

digital corpora used to train Large Language Models (LLMs). Consequently, most existing Arabic LMs are predominantly MSA-centric, with limited and disproportionate coverage of diverse dialects (Antoun et al., 2020; Inoue et al., 2021; Abdulmageed et al., 2021; Antoun et al., 2021; Sengupta et al., 2023; Lan et al., 2020). While these models are often implicitly expected to generalize across the Arabic continuum, cross-lingual transfer is not guaranteed, particularly as many dialects exhibit low mutual intelligibility with MSA and one another (Abu Farha and Magdy, 2022; Kegel et al., 2023). This linguistic imbalance has spurred the development of dialect-specific models (Qarah, 2024b; Gaanoun et al., 2025; Abdaoui et al., 2021; Qarah, 2024a; Shang et al., 2025; AlYami and Al-Zaidy, 2022). However, the necessity of such specialization remains an open question: **Do MSA-centric or multi-dialect models transfer equitably across the Arabic dialectal landscape?** To investigate this, we probe the internal representations of Arabic LMs to evaluate their transferability to various dialects across three core NLP tasks—Sentiment Analysis (SA), Named Entity Recognition (NER), and Part-of-Speech (POS) tagging—benchmarking their performance against specialized, dialect-specific counterparts.

While probing provides a functional account of cross-lingual transfer, its scope is often constrained by specific task selections and potential confounding dataset characteristics. Prior research has leveraged annotated linguistic features to predict inter-lingual transfer or interference (Lin et al., 2019; Eronen et al., 2023); however, the scarcity of granular linguistic annotations for Arabic dialects renders such approaches currently unfeasible. While Alsudais et al. (2022) utilized lexical overlap within the parallel MADAR corpus (Bouamor et al., 2018) as a proxy for dialectal similarity, surface-level overlap fails to adequately capture deeper syntactic and semantic nuances. Given the proficiency of

¹Our code: https://github.com/muizzkhalak/cross_lingual_transfer_arabic

LMs in encoding these high-level features (Conneau et al., 2018a), we propose employing Representational Similarity Analysis (RSA) between MSA and DA model representations. Specifically, we utilize Centered Kernel Alignment (CKA) (Kornblith et al., 2019) to quantify layer-wise similarity between DA and MSA models using parallel sentences from the MADAR dataset. We posit that this provides an intrinsic metric of cross-lingual transfer: high representational similarity between MSA and DA models serves as a robust signal for effective transferability and suggests that inter-dialectal interference will be minimized.

We integrate these two methodologies—Probing and Representational Similarity Analysis (RSA)—to provide a holistic evaluation of cross-lingual transfer between MSA and its dialects. Our findings indicate that, in general, MSA-centric models demonstrate strong transferability to DA, occasionally outperforming dialect-specific models. However, this transfer remains significantly disproportionate across the dialectal spectrum. To investigate the drivers of this disparity, we correlate model performance with factors such as geographic proximity and pretraining data volume. Our analysis reveals that dialect-specific models consistently exceed the performance of MSA and multi-dialect models only when supported by substantial DA pretraining data. Furthermore, geographic proximity serves as a strong predictor of transferability, aligning with the dialectal continuum hypothesis: as geographic distance increases, mutual intelligibility diminishes, posing a greater challenge for MSA-to-DA transfer.

2 Related Work

Cross-Lingual Transfer. Prior work consistently shows that transfer effectiveness depends strongly on source–target similarity: typologically or linguistically closer languages yield better zero-shot transfer, and language–distance metrics correlate with downstream accuracy (Eronen et al., 2023; Philippy et al., 2023). Consequently, selecting a related transfer source—rather than defaulting to English—can substantially improve performance for low-resource targets (Lin et al., 2019; Eronen et al., 2023). Multilingual LMs such as mBERT further demonstrate robust cross-lingual generalization without explicit alignment, achieving competitive zero-shot results across tasks like Natural Language Inference (NLI) and NER after fine-

tuning on a single language (Conneau et al., 2018b; Pires et al., 2019; Wu and Dredze, 2019). This behavior is often attributed to increasingly language-agnostic representations in higher layers (K et al., 2020; Artetxe et al., 2020); notably, retraining only the embedding layer of a monolingual transformer can effectively adapt it to a new language, suggesting that upper-layer features are largely transferable (Artetxe et al., 2020). However, multilingual training often leads to negative interference especially for high resource languages (Bapna and Firat, 2019; Chang et al., 2024; Alastruey et al., 2025). Overall, these findings motivate balanced transfer strategies, where leveraging one or a few related high-resource languages yields strong gains for low-resource targets without incurring excessive interference (Seto et al., 2025). In this study, we examine how well MSA-centric and multi-dialect Arabic models transfer to DA varieties, and whether negative interference occurs in dialectal context, even when linguistic similarity is assumed. Although our process follows cross-dialectal transfer across Arabic models, we use "cross-lingual" in the rest of the paper to align with literature norms.

Probing Arabic Language Models. Previous work has examined how linguistic information is distributed across layers and neurons in Arabic transformer models, including dialect-trained variants (Abdelali et al., 2022). This line of research reports that lower and middle layers primarily encode morphology, upper layers capture syntactic dependencies, and MSA-based models often fail to represent dialect-specific nuances despite substantial vocabulary overlap; embedding-layer neurons tend to be polysemous, whereas mid-layer neurons are more specialized for particular linguistic properties (Abdelali et al., 2022). By contrast, we focus on model-level representations: we assess how MSA, general-purpose Arabic (MIX/Multi-DA), and dialect-specific encoders encode morphological and syntactic information via controlled probing tasks, comparing best-layer performance across models rather than inspecting individual neurons.

Similarity of Arabic Dialects. Recent work in Arabic dialectometry links textual similarity to geography, showing that geographically proximate varieties tend to be linguistically closer, sometimes more so than political boundaries (Alsudais et al., 2022). Across dialectal corpora, classic NLP/IR similarity measures further reveal strong Levantine cohesion and suggest that Palestinian is par-

ticularly close to MSA (Kwaik et al., 2018), while comparisons with Classical Arabic provide a historical reference for modern variation (Abouzahr and Abouzahr, 2025). Analyses of the MADAR corpus also report Muscat (Spoken in Oman) as among the closest dialects to MSA and Sfax (Spoken in Tunisia) among the most distant (Bouamor et al., 2018). These lexical analyses remain limited at capturing syntactic and semantic similarity, therefore, our work studies whether such proximity relationships are internalized by pre-trained language models, which were shown to capture syntactic and semantic features.

3 Methodology

In this section, we outline our methodological framework for studying cross-lingual transfer of Arabic LMs. We combine two complementary methods for this study, namely: linear probes trained on frozen layer-wise embeddings (Section 3.1), and similarity of model representations measured using Centered Kernel Alignment (CKA) (Section 3.2). Furthermore, motivated by the linguistic theory of dialect continuum and its implications (Chambers and Trudgill, 1998), we propose a proxy for measuring geographic proximity to MSA (Section 3.3).

3.1 Probing Classifiers

Probing has been widely applied to analyze multilingual models, where it can reveal whether models encode language-specific features or universal cross-linguistic structures (Tenney et al., 2019b). In the case of Arabic, probing classifiers can help determine whether dialectal and MSA models capture similar linguistic phenomena and whether knowledge acquired in one variety (e.g., MSA) transfers to another (e.g., Egyptian or Levantine Arabic). By training probes on embeddings derived from dataset of one variety using a language model of different variety.

The standard probing setup involves training a lightweight classifier—commonly a linear or shallow feedforward model—on top of frozen representations extracted from different layers of the LM. We train probes to predict a supervised linguistic property (e.g., POS tags) using embeddings as input. If the probe achieves high accuracy, this suggests that the relevant linguistic information is encoded in the representations at that layer and is easily extractable by a simple function (Hewitt

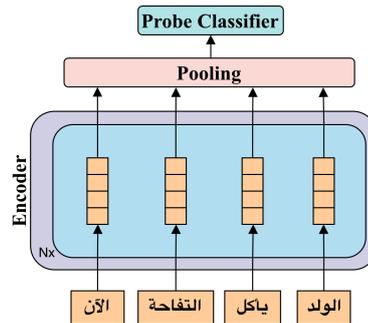


Figure 1: Architecture of the probing classifier for the example sentence “The boy is eating the apple now.” Sentence representations pass through N layers, and each layer is probed using the classifier in Eq. 1.

and Manning, 2019). Probing classifiers are deliberately kept simple to minimize the chance that the probe itself learns the task from scratch. Instead, their purpose is to act as a diagnostic tool that reveals the presence (or absence) of linguistic features in the LM representations (Pimentel et al., 2020). Figure 1 illustrates the probing architecture, where we employ a linear classifier with a softmax activation layer. For a given layer l , the probability distribution y_l is computed as:

$$y_l = \text{softmax}(Wz_l + b) \quad (1)$$

where z_l is a representation derived from the hidden states $H_l = \{h_{l,1}, \dots, h_{l,T}\}$. We define two pooling strategies depending on the probing task:

- **Token-level** to POS and NER tasks: $z_l = h_{l,i}$, the embedding at index i .
- **Sentence-level** for the SA task: $z_l = \frac{1}{T} \sum_{i=1}^T h_{k,i}$, the mean-pooled average of all token embeddings in the sequence.

3.2 Representation Similarity Analysis

Besides cross-lingual transfer as measured using probing on selected tasks, we also compute representation similarity, which can be viewed as an intrinsic signal for transfer. Primarily, we seek to study whether MSA LMs learn similar representations to their DA counterpart, which can be seen as a signal of similarity between DA varieties and MSA, and of the sufficiency of MSA for capturing the nuances of DA varieties.

Figure 2 illustrates our architecture for representation similarity, where we use Centered Kernel Alignment (CKA) (Kornblith et al., 2019) to measure the similarity of hidden representations. Let $X \in \mathbb{R}^{n \times d_x}$ and $Y \in \mathbb{R}^{n \times d_y}$ be representations of the

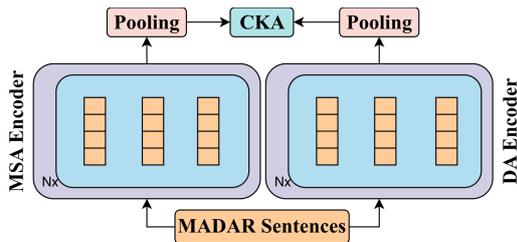


Figure 2: Architecture of CKA for representation similarity. MADAR parallel sentences are encoded by MSA and DA encoders through N layers, and the resulting representations are compared using linear CKA (Eq. 2).

same n inputs. For linear CKA, we set $K = XX^\top$ and $L = YY^\top$ and compute:

$$\text{CKA}(X, Y) = \frac{\|X^\top Y\|_F^2}{\|X^\top X\|_F \|Y^\top Y\|_F} \in [0, 1], \quad (2)$$

where higher values indicate the representations are similar. Linear CKA is invariant to orthogonal transformations and isotropic rescaling of features, making it robust to rotations and global scaling differences between models.

This yields a compact and task-agnostic view of how similarly MSA and dialectal models encode Arabic input, complementing our probing-based analysis of functional transfer.

3.3 Evaluating Dialectal Continuum

We test whether similarity measures track geographical proximity continuum, i.e., whether models associated with countries closer (geographically) to a chosen MSA anchor exhibit higher similarity to an MSA-pretrained encoder. Similar approaches have been used for finding a relationship between language distance and cross-lingual transfer (Eronen et al., 2023; Philippou et al., 2023). A key challenge is that MSA has no fixed geographic locus, as it is used in formal contexts across all Arab countries. The literature proposes several proxies: one strand argues that MSA is particularly close to Palestinian (Levantine) Arabic, based on lexical similarity within Levantine and North-African varieties (Kwaik et al., 2018), however, this study did not include Gulf dialects. Other work highlights the strong affinity between Classical Arabic (CA) and Yemeni Arabic, citing the preservation of archaic forms, phonological features, and relative isolation, and notes that Gulf Arabic, despite borrowings (e.g., from Persian), maintains close ties to Bedouin dialects (Boyi et al., 2024). A separate lexical / lexical-semantic comparison

with CA across regions also identifies Yemeni Arabic as the closest modern variety (Abouzahr and Abouzahr, 2025); although the authors attribute this in part to corpus-size bias for Yemeni data, they also show that Gulf Arabic—geographically close to Yemen—has high similarity to CA. Finally, analysis of the MADAR corpus reports that Muscat (Oman) is the dialect most similar to MSA (Bouamor et al., 2018), and Muscat is likewise geographically near Yemen. Given the widely attested closeness of MSA to CA, we therefore **use Yemen as a geographic proxy for the MSA anchor**, while explicitly acknowledging this as an operational choice.

4 Experimental Setup

4.1 Probing Datasets

For fair cross-dialect comparison, we select dialectal datasets per task and enforce equal sample sizes per probe: for **POS/NER** we balance by number of labeled tokens, and for **SA** by number of sentences. General-purpose Arabic models are consistently evaluated against dialect-specific models on the latter’s native test sets. Table 2 details the corpora utilized for each dialect; for multi-dialectal datasets, we partition and reuse relevant subsets for their corresponding varieties. Due to the absence of publicly available NER corpora for Gulf dialects (specifically Riyadh and Muscat), we derive a silver-standard dataset from the MADAR corpus (Bouamor et al., 2018). We extract 2,000 sentences per dialect and perform automated entity annotation using the **CAMEL-Lab/bert-base-arabic-camelbert-da-ner** model (Inoue et al., 2021).

For the POS tagging task, we apply preprocessing tailored to Arabic cliticization, where multiple grammatical units are written as a single orthographic word. For example, **وموبايل** (“**and mobile**”) combines **و** (*CONJ*) and **موبايل** (*NOUN*); we normalize such cases by aggregating clitic and stem tags into composite labels (e.g., *CONJ+NOUN*), following prior work in Arabic POS tagging (Darwish et al., 2018; Habash and Rambow, 2006). NER datasets are used with their original tags and no additional preprocessing. For sentiment analysis, where data often comes from tweets and other social media, we remove URLs, HTML tags, usernames, and strip emojis using the emoji² library.

²<https://pypi.org/project/emoji/>

Dialect identification filtering. Because many probing datasets originate from noisy online sources, they may contain *off-target* sentences (other dialects or MSA). To mitigate this, we run automatic dialect identification on every sentence using **CAMeL-Lab/bert-base-arabic-camelbert-mix-did-madar-corpus26** (Inoue et al., 2021). City-level predictions (e.g., *Riyadh, Jeddah, Cairo*) are mapped to countries, and we retain only sentences whose predicted country matches the dataset’s target dialect (or reassign them when the dataset explicitly spans multiple countries). This filtering step yields cleaner, better-aligned inputs for subsequent probing analyses.

4.2 Similarity Analysis Dataset

Our similarity analysis requires access to parallel sentences between MSA and dialects. Therefore, we use the **MADAR** corpus (Bouamor et al., 2018), which provides aligned versions of the same sentences in MSA and 25 city-level dialects. The corpus contains 2000 parallel sentences per city which helps avoid sample-size artifacts.

4.3 Models

We focus on BERT-style encoders throughout this study, so that cross-model comparisons are not confounded by architectural differences and transfer effects remain directly comparable. Table 1 summarizes the models used and their key properties that help contextualize their behavior and performance across our experiments.

4.4 Probing

Probe Training We employ a linear statistical classifier—multinomial logistic regression (softmax over a single affine layer)—implemented with the NEUROX toolkit (Dalvi et al., 2019). For the word-level tasks (POS and NER), we aggregate subword representations by selecting the embedding of the final subword token (Liu et al., 2019). For sentence-level Sentiment Analysis, representations are derived via mean-pooling across all tokens in the sequence. We use stratified 80/20 train/test splits, the Adam optimizer with a small learning rate. Each layer ℓ is probed independently, producing a performance curve per model and task. Token-level embeddings for POS/NER are obtained by aggregating subword representations to the word level (using the embedding of the last subword token (Liu et al., 2019)), while sentence-level embeddings for sentiment are computed by mean-pooling

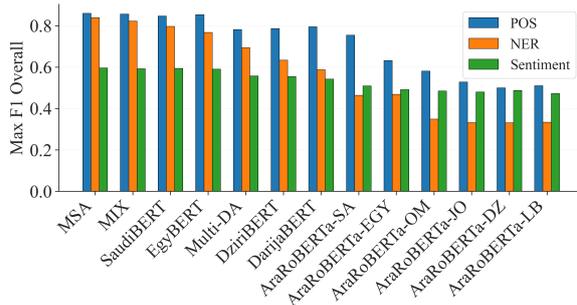


Figure 3: Performance of best performing layer on MSA Tasks.

token representations over the sequence.

Probe Evaluation We use the F1 score for the evaluation. We measure the performance of General models (i.e. MSA, Multi-DA and MIX) and DA models on both MSA only and DA datasets. We think that this offers a bidirectional look at the cross-lingual transfer of Arabic models. Where DA model that can transfer to MSA might indicate that the dialect is close to MSA and vice versa. Unless specified differently, we report the best F1 score across the model layers.

4.5 Representation Similarity

We use CKA to measure the representation similarity between MSA and DA models. We compute hidden states on the MADAR dataset, take the mean over the sentence tokens, then measure the layer-wise CKA similarity. The fact that MADAR is a parallel dataset allows us to investigate the following three scenarios:

1. MSA model vs. Dialect model representation of DA sentences.
2. MSA model vs. Dialect model representation of MSA sentences.
3. MSA model vs. Dialect model representation of MSA and DA sentences respectively.

Finally, since MADAR contains city level dialectal sentences and some countries are represented by more than one city (e.g., Egypt: Alexandria, Aswan, Cairo), we take the mean of CKA similarity over the cities to obtain the similarity per country, and consequently, per dialect.

Model Family	Model Name	Training Varieties	Tokens
CAMeLBERT	CAMeLBERT-MSA (Inoue et al., 2021)	MSA	12.6B
	CAMeLBERT-DA (Inoue et al., 2021)	DA	5.8B
	CAMeLBERT-MIX (Inoue et al., 2021)	MSA + DA + CA	17.3B
Dialect-Specific	SaudiBERT (Qarah, 2024b)	Saudi Dialect	2.7B
	EgyBERT (Qarah, 2024a)	Egyptian Dialect	1.07B
	DarijaBERT (Gaanoun et al., 2025)	Moroccan Dialect (Darija)	100M
	DziriBERT (Abdaoui et al., 2021)	Algerian Dialect	20M
AraRoBERTa	AraRoBERTa-SA (AlYami and Al-Zaidy, 2022)	Saudi Dialect	45.4M
	AraRoBERTa-EGY (AlYami and Al-Zaidy, 2022)	Egyptian Dialect	37.2M
	AraRoBERTa-OM (AlYami and Al-Zaidy, 2022)	Omani Dialect	3.8M
	AraRoBERTa-LB (AlYami and Al-Zaidy, 2022)	Lebanese Dialect	3.6M
	AraRoBERTa-JO (AlYami and Al-Zaidy, 2022)	Jordanian Dialect	2.6M
	AraRoBERTa-DZ (AlYami and Al-Zaidy, 2022)	Algerian Dialect	1.9M

Table 1: Overview of Arabic Pre-trained Language Models (PLMs)

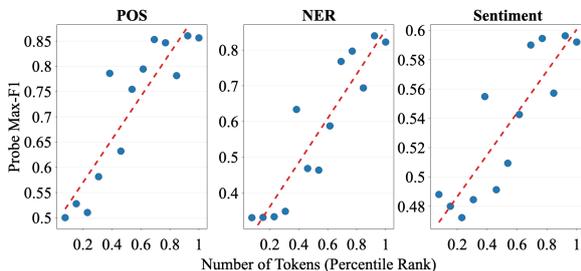


Figure 4: Impact of pretraining corpus size on probe performance across tasks. The Percentile Rank of the number of tokens is displayed for better visual interpretation.

5 Results and Analysis

5.1 Probing Analysis

For probing experiments implemented in Section 4.4, we analyze and compare the performance of the best layer for each model. This is based on evidence that models store different information disproportionately in different layers (Tenney et al., 2019a), while we are interested in comparing models in terms of their absolute performance rather than per-layer comparison.

Figure 3 reports the macro-F1 of the best-performing layer for each model and task on the MSA datasets, as discussed in Section 4.4. Across POS, NER, and SA, a consistent pattern emerges: the **MSA-pretrained** model achieves the strongest performance on MSA text, outperforming both the **MIX** and **Multi-DA** models. Among dialect-specific encoders, **Saudi** and **Egyptian** models typically rank closest to the MSA model, while other regional models lag more substantially. The same trend holds within the AraRoBERTa family, where Saudi and Egyptian variants systematically outperform other dialectal variants on MSA inputs. By

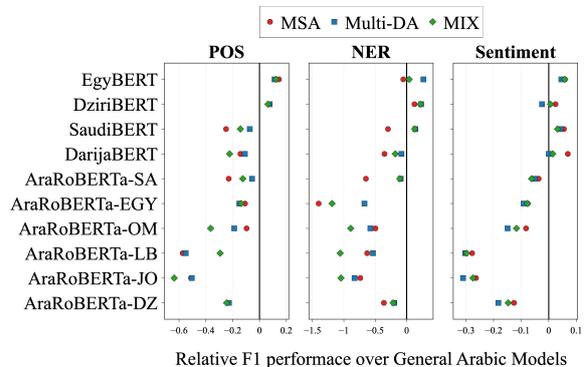


Figure 5: Relative performance of general vs. dialect-specific Arabic models on native dialectal datasets. Points to the right of the reference line denote cases where the dialect-specific model achieves higher performance, while points to the left indicate that the general model remains superior.³

contrast, models trained on North African varieties, such as DziriBERT and DarijaBERT, generalize poorly to MSA across all tasks, likely reflecting the pronounced lexical and morphological divergence between Maghrebi dialects and MSA. When we relate model training size, as summarized in Table 1, to performance on MSA data, we observe the trend visualized in Figure 4: models trained on larger and more diverse corpora tend to exhibit stronger generalization, underscoring the importance of corpus scale and linguistic variety for cross-dialectal robustness.

Figure 5 compares an **MSA** model, a **Multi-DA** model, and a **MIX** model against DA encoders on dialectal POS, NER, and SA. We compare models at their best-performing layer to control for layer specialization effects. We report the macro-

³Non-normalized F1 score results are provided in our code repository.

F1 score performance of each model relative to the dialect specific model. We compute the relative difference as follows:

$$\Delta = \frac{F1_{DA} - F1_M}{F1_{DA}} \quad (3)$$

where $M = \{\text{MSA, Multi-DA, MIX}\}$. We normalize by $F1_{DA}$ to make the results comparable across dialects.

Overall, DA models such as EgyBERT and DziriBERT consistently excel on their native dialects—especially for SA and for POS where morphology and cliticization diverge most from MSA—highlighting the benefits of dialect-specific pretraining. The MSA model, while weaker on SA for most dialects, remains competitive on structural tasks (POS, NER) for several varieties, indicating that MSA pretraining preserves transferable syntactic and entity-structure regularities. The Multi-DA model provides a strong dialect-agnostic baseline, where it often outperforms the MSA model on the SA task, but often falls short of the strongest DA models on their home dialect although it is trained on more data. This suggests that training on multiple Arabic dialects can still lead to negative interference (Wang et al., 2020), although they are regarded as similar in literature. The MIX model shows a similar breadth-sharpness trade-off: it can match or surpass many dialectal models on POS, but tends to underperform DA encoders on NER and SA when domain and style are highly dialect-specific. Across these comparisons, pretraining scale emerges as a key factor. DA models that reliably beat the general-purpose encoders (MSA, Multi-DA, MIX) are often trained on substantially larger pretraining corpora than lighter dialectal models (e.g., EgyBERT vs. AraRoBERTa-EGY, DziriBERT vs. AraRoBERTa-DZ). This suggests that the amount of pretraining data is crucial to improve the model’s encoding of dialect-specific morphological, syntactic, and semantic features.

5.2 Representation Similarity Analysis

Figure 6 summarizes the average layer-wise CKA between the MSA model and each DA model across all three scenarios discussed in Section 4.5. When both models encode the same dialectal sentences (Scenario 1), the highest alignment is observed for SaudiBERT on Saudi text, followed by DziriBERT, EgyBERT, and DarijaBERT, with AraRoBERTa-LB showing the closest alignment among the AraRoBERTa variants. When both

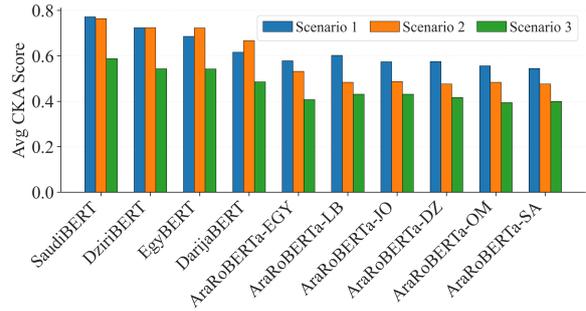


Figure 6: Average layer-wise CKA between the MSA model and DA models for the three scenarios described in Section 4.5.

encoders process the MSA versions of the same sentences (Scenario 2), the order of BERT-based dialect models remains similar, while within the AraRoBERTa family, AraRoBERTa-EGY becomes the closest to the MSA encoder and the remaining variants form a slightly lower, tightly clustered group. In scenario 3, where the MSA model encodes MSA sentences and dialectal models encode parallel dialect sentences, the BERT models ranking largely persists (SaudiBERT > DziriBERT > EgyBERT > DarijaBERT), whereas within AraRoBERTa, Levantine variants (AraRoBERTa-JO, AraRoBERTa-LB) show the strongest alignment, followed by North-African variants (AraRoBERTa-DZ, AraRoBERTa-EGY), then Gulf variants (AraRoBERTa-SA, AraRoBERTa-OM). Across scenarios, scenario 3 has the lowest similarity scores, showing the representations of DA and MSA models are not agnostic to their input variant. Furthermore, even when the input is the same variant (i.e. scenario 1 and 2), the similarity is always lower than 0.8, showing that MSA models fail to capture DA specific nuances and vice versa. The results of AraRoBERTa models suggest that this can be partly explained by the amount of training data, where the similarity of these models is lower than the BERT-based models, which are trained on more data.

5.3 Proximity Analysis

As discussed in Sections 5.2 and 5.1, there is significant disparity between dialect specific models in terms of their probing performance and similarity with MSA models. Motivated by the documented linguistic phenomena of dialect continuum (Chambers and Trudgill, 1998), and by previous studies that investigated the relationship between geographic proximity and lexical overlap (Alsu-

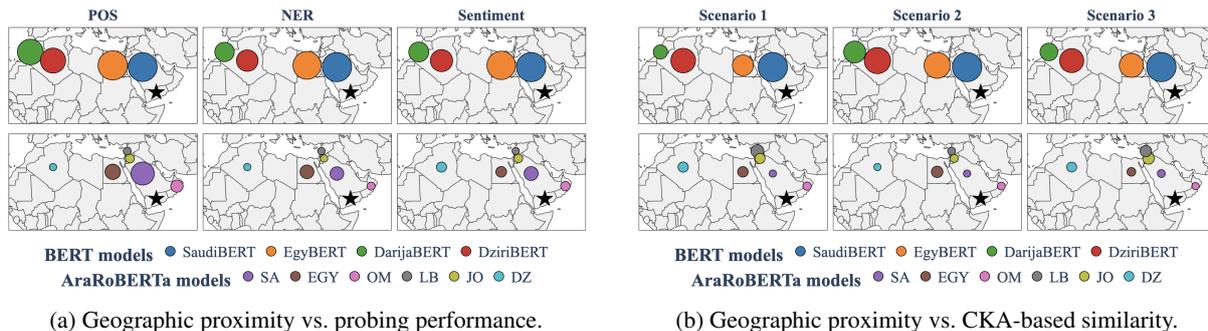


Figure 7: Comparison of Geographic Proximity with Probing Performance and CKA Similarity. The black star marks Yemen (the MSA anchor), and each bubble denotes a DA model; bubble size encodes (a) macro-F1 score on MSA and (b) average CKA score.

dais et al., 2022), we study whether geographic proximity can explain the disparity in probing and representation similarity results. Note that we use Yemen as the geographic location for MSA as explained in Section 3.3. Figure 7a shows dialectal models by their geographic location, with bubble sizes proportional to their macro-F1 score on MSA text from the probing tasks. Across all tasks, we observe a clear dialectal continuum: models associated with regions geographically closer to the MSA anchor are represented by larger bubbles, which gradually shrink with increasing distance. This pattern indicates that cross-lingual transfer in terms of *functional* similarity aligns well with geographic proximity. Figure 7b presents an analogous view for CKA-based similarity. Here, the dialectal continuum is evident primarily for models trained on large-scale corpora, while models with more limited pretraining data (i.e. AraRoBERTa models) do not reflect a dialect continuum effect. The results suggest that proximity plays a role in model performance and similarity, and is in line with the phenomena of dialect continuum, where more distant dialects and less mutually intelligible, which leads to weaker cross-lingual transfer and lower similarity. Specifically, Scenario 2 in Figure 7b is more comparable with Figure 7a since we compare the dialectal continuum of the MSA model and the DA models on MSA text.

6 Discussion

Based on the probing results on MSA datasets, we observe that the MSA model outperforms both MIX and Multi-DA models. The relative underperformance of the MIX model aligns with prior findings that monolingual MSA encoders surpass mixed-variant models on MSA tasks (Inoue et al., 2021).

Although the MIX model is trained on MSA data, the underperformance can be due to curse of multilinguality which adversely affects high-resource varieties (Chang et al., 2024). Among the DA models we observe that cross-lingual transfer to MSA is strongly dependent on the amount of data the model is trained on. This observation is further validated by geographical proximity analysis where models trained on large amount of data capture dialect continuum more efficiently. We also observe that it is easier for the language models to capture functional similarity even with less training data. But it is more difficult for Arabic language models to capture structural similarity as observed from geographical proximity analysis. Furthermore, probing and CKA reveal a clear functional–structural gap: some models (e.g., DziriBERT) can be structurally close to MSA yet fail to convert this into superior probing performance. This is in line with prior work showing that high CKA similarity does not guarantee functionally similar or equally useful features (Davari et al., 2023).

The probing performance on the DA datasets shows that the MSA-pretrained encoder is generally stronger on NER and POS tasks, while DA models tend to excel on SA tasks. The performance on POS suggests that Arabic and its dialects maintain a similar syntax, while the performance on NER can be attributed to code switching, where DA varieties often use MSA for entities. However, this does not translate to a higher performance on SA, where higher level semantic and pragmatic understanding is required to predict the sentiment. These findings suggest that Arabic dialects are similar to MSA in their syntax and use code switching for entities which explains their performance on POS and NER tasks respectively, however, they struggle on SA which is characterized by dialect

specific idioms, non-standard orthography and distinct negation patterns that are absent from their pretraining data.

Comparing DA models with *general* dialectal models (Multi-DA; MIX) on downstream tasks (Figures 5) shows that *general* models frequently achieve stronger average performance across dialects, a likely consequence of larger and more diverse pretraining. The MIX model commonly surpasses the Multi-DA model, suggesting that incorporating MSA and Classical Arabic alongside dialectal text improves the encoding of morphological and syntactic regularities. Nonetheless, there are notable exceptions: DziribERT and EgyBERT often outperform *general* models on their native dialects. For SA, multiple DA models (with the exception of some AraRoBERTa variants) outperform general models, while the MIX model retains advantages on NER and POS tasks, where training on MSA text gives an advantage. We see that for most high resource dialects (Egyptian, Saudi Arabia), DA models usually outperform *general* dialect models. Whereas for low resource dialects (Lebanese, Jordan), *general* dialect models outperform native dialectal models. This observation is consistent with recent research on curse of multilinguality (Chang et al., 2024; Wu and Dredze, 2020), where multi-dialect pretraining benefits low resource dialects, but negatively affects high resource dialect performance.

7 Conclusion

In this work, we study how Arabic language models (LMs) capture cross-lingual transfer among dialects and their relationship to Modern Standard Arabic (MSA), combining **functional** evidence from probing classifiers (POS, NER, SA) and **structural** evidence from representation similarity (CKA), and linking both to **geographic proximity**. Our analysis covered MSA, MIX, and Multi-DA encoders alongside mono-dialectal DA models. We find consistent signals of a dialectal continuum: models associated with geographically closer varieties tend to align more closely with MSA in both probing and CKA, but this effect strongly depends on the amount of pretraining data. Interestingly, even among dialects of the same language we find evidence for negative interference, where Multi-DA models underperform mono-dialectal models of high resource dialects. Our findings question the assumption of cross-lingual transfer from MSA to

its dialects, where transfer is feasible but disproportionate across varieties. Furthermore, Multi-DA models might benefit from dialect or dialect-region specific parameters to avoid negative interference.

Limitations

Our work has the following limitations:

Dialectal dataset coverage: No publicly available Gulf-dialect NER dataset exists to our knowledge; we therefore leveraged MADAR and produced proxy annotations using an existing NER model. This introduces potential annotation noise and domain mismatch that may affect NER-related comparisons. Additionally, we rely on regional datasets (e.g., Gulf, Maghrebi) that are further segmented by country (e.g., Saudi, Moroccan) using automatic dialect identification. However, dialect identification is imperfect, and occasional false positives may introduce noise into the country-level splits and, consequently, into our downstream analyses.

Geographic proxy for MSA: Our choice of Yemen as a geographic proxy for MSA is literature-informed but not definitive. Results involving distance-to-anchor should be interpreted as operational rather than canonical; alternative anchors could yield different effect sizes.

Model heterogeneity: The compared models differ in training size, vocabularies, and pretraining hyperparameters. Such heterogeneity can confound attribution of observed effects. A more controlled comparison—equalizing corpus size, and training settings—would yield more robust findings.

Pretraining data composition: The paper does not focus the effect of the composition of pretraining data of dialectal models. Some models, such as DarijaBERT, were trained by filtering out MSA data, whereas SaudiBERT, DziribERT, and EgyBERT were trained on data collected based on geolocation, which may include traces of MSA data. The presence of noisy MSA data in dialectal models may affect the model’s performance.

References

- Amine Abdaoui, Mohamed Berrimi, Mourad Oussalah, and Abdelouahab Moussaoui. 2021. [Dziribert: a pre-trained language model for the algerian dialect](#). *CoRR*, abs/2109.12346.
- Ahmed Abdelali, Nadir Durrani, Fahim Dalvi, and Hassan Sajjad. 2022. [Post-hoc analysis of Arabic transformer models](#). In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting*

- Neural Networks for NLP*, pages 91–103, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Muhammad Abdul-Mageed, AbdelRahim Elmadany, and El Moatez Billah Nagoudi. 2021. **ARBERT & MARBERT: Deep bidirectional transformers for Arabic**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7088–7105, Online. Association for Computational Linguistics.
- Hossameddine Abouzahr and Mohammed Abouzahr. 2025. **Lexical and lexical-semantic comparisons of classical arabic and dialects**. *International Journal of Arabic Linguistics*, 11:66–98.
- Ibrahim Abu Farha and Walid Magdy. 2022. **The effect of Arabic dialect familiarity on data annotation**. In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 399–408, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Kathrein Abu Kwaik, Motaz Saad, Stergios Chatzikiriakidis, and Simon Dobnik. 2018. **Shami: A corpus of Levantine Arabic dialects**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ashraf Ahmad, Mohammad Azzeh, Eman Alnagi, Qasem Abu Al-Haija, Dana Halabi, Abdullah Aref, and Yousef AbuHour. 2024. **Hate speech detection in the arabic language: corpus design, construction, and evaluation**. *Frontiers in Artificial Intelligence*, Volume 7 - 2024.
- Belen Alastruey, João Maria Janeiro, Alexandre Al-lauzen, Maha Elbayad, Loïc Barrault, and Marta R. Costa-jussà. 2025. **Interference matrix: Quantifying cross-lingual interference in transformer encoders**. *Preprint*, arXiv:2508.02256.
- Mashaël AlDuwais, Hend Al-Khalifa, and Abdulmalik AlSalman. 2024. **CLEANANERCorp: Identifying and correcting incorrect labels in the ANERcorp dataset**. In *Proceedings of the 6th Workshop on Open-Source Arabic Corpora and Processing Tools (OSACT) with Shared Tasks on Arabic LLMs Hallucination and Dialect to MSA Machine Translation @ LREC-COLING 2024*, pages 13–19, Torino, Italia. ELRA and ICCL.
- Areeb Alowisheq, Nora Al-Twairish, Mawaheb Al-tuwaijri, Afnan Almoammar, Alhanouf Alsuwaillem, Tarfa Albuhairei, Wejdan Alahaideb, and Sarah Alhummoud. 2021. **Marsa: Multi-domain arabic resources for sentiment analysis**. *IEEE Access*, 9:142718–142728.
- Abdulkareem Alsudais, Wafa Alotaibi, and Faye Alomary. 2022. **Similarities between arabic dialects: Investigating geographical proximity**. *Information Processing & Management*, 59(1):102770.
- Mohamed Aly and Amir Atiya. 2013. **LABR: A large scale Arabic book reviews dataset**. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 494–498, Sofia, Bulgaria. Association for Computational Linguistics.
- Reem AlYami and Rabah Al-Zaidy. 2022. **Weakly and semi-supervised learning for Arabic text classification using monodialectal language models**. In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 260–272, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. **AraBERT: Transformer-based model for Arabic language understanding**. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15, Marseille, France. European Language Resource Association.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2021. **AraGPT2: Pre-trained transformer for Arabic language generation**. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pages 196–207, Kyiv, Ukraine (Virtual). Association for Computational Linguistics.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. **On the cross-lingual transferability of monolingual representations**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Ankur Bapna and Orhan Firat. 2019. **Simple, scalable adaptation for neural machine translation**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Zakaria Benmounah, Abdennour Boulesnane, Abdeladim Fadheli, and Mustapha Khial. 2023. **Sentiment analysis on algerian dialect with transformers**. *Applied Sciences*, 13(20).
- Houda Bouamor, Nizar Habash, Mohammad Salameh, Wajdi Zaghoulani, Owen Rambow, Dana Abdulrahim, Ossama Obeid, Salam Khalifa, Fadhil Eryani, Alexander Erdmann, and Kemal Oflazer. 2018. **The MADAR Arabic dialect corpus and lexicon**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Ahmad Muhammad Boyi, Muhammad Badamasi Yusuf, and Mustapha Muhammad Isa. 2024. **The relevance**

- of arabic dialects and their diversity across the arab world. *Journal of Arts and Sociological Research*, 6.
- J. K. Chambers and Peter Trudgill. 1998. *Dialectology*, page 3–12. Cambridge Textbooks in Linguistics. Cambridge University Press.
- Tyler A. Chang, Catherine Arnett, Zhuowen Tu, and Ben Bergen. 2024. [When is multilinguality a curse? language modeling for 250 high- and low-resource languages](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4074–4096, Miami, Florida, USA. Association for Computational Linguistics.
- Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. 2018a. [What you can cram into a single \\$&!#* vector: Probing sentence embeddings for linguistic properties](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2126–2136, Melbourne, Australia. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018b. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Abdelhalim Hafedh Dahou and Mohamed Amine Cheragui. 2023. [Dzner: A large algerian named entity recognition dataset](#). *Natural Language Processing Journal*, 3:100005.
- Fahim Dalvi, Avery Nortonsmith, Anthony Bau, Yonatan Belinkov, Hassan Sajjad, Nadir Durrani, and James Glass. 2019. [Neurox: A toolkit for analyzing individual neurons in neural networks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):9851–9852.
- Kareem Darwish, Hamdy Mubarak, Ahmed Abdelali, Mohamed Eldesouki, Younes Samih, Randah Alharbi, Mohammed Attia, Walid Magdy, and Laura Kallmeyer. 2018. [Multi-dialect Arabic POS tagging: A CRF approach](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- MohammadReza Davari, Stefan Horoi, Amine Natic, Guillaume Lajoie, Guy Wolf, and Eugene Belilovsky. 2023. [Reliability of CKA as a similarity measure in deep learning](#). In *The Eleventh International Conference on Learning Representations*.
- Niama Elkhbir, Urchade Zaratiana, Nadi Tomeh, and Thierry Charnois. 2023. [Cross-dialectal named entity recognition in Arabic](#). In *Proceedings of ArabicNLP 2023*, pages 140–149, Singapore (Hybrid). Association for Computational Linguistics.
- Juuso Eronen, Michal Ptaszynski, and Fumito Masui. 2023. [Zero-shot cross-lingual transfer language selection using linguistic similarity](#). *Information Processing & Management*, 60(3):103250.
- Charles A. Ferguson. 1959. [Diglossia](#). *WORD*, 15(2):325–340.
- Kamel Gaanoun, Abdou Mohamed Naira, Anass Al-lak, and Imade Benellallam. 2025. [Darijabert: a step forward in nlp for the written moroccan dialect](#). *International Journal of Data Science and Analytics*, 20(2):917–929.
- Nizar Habash, Muhammed AbuOdeh, Dima Taji, Reem Faraj, Jamila El Gizuli, and Omar Kallas. 2022. [Camel treebank: An open multi-genre Arabic dependency treebank](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2672–2681, Marseille, France. European Language Resources Association.
- Nizar Habash and Owen Rambow. 2006. [MAGEAD: A morphological analyzer and generator for the Arabic dialects](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 681–688, Sydney, Australia. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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*, pages 92–104, Kyiv, Ukraine (Virtual). Association for Computational Linguistics.
- Mustafa Jarrar, Mohammed Khalilia, and Sana Ghanem. 2022. [Wojood: Nested Arabic named entity corpus and recognition using BERT](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3626–3636, Marseille, France. European Language Resources Association.
- Mouad Jbel, Mourad Jabrane, Imad Hafidi, and Abdulmutallib Mtrane. 2024. [Sentiment analysis dataset in moroccan dialect: bridging the gap between arabic and latin scripted dialect](#). *Lang. Resour. Eval.*, 59(2):1401–1430.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual bert: An empirical study](#). In *International Conference on Learning Representations*.

- Amr Keleg, Sharon Goldwater, and Walid Magdy. 2023. [ALDi: Quantifying the Arabic level of dialectness of text](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10597–10611, Singapore. Association for Computational Linguistics.
- Salam Khalifa, Nizar Habash, Dana Abdulrahim, and Sara Hassan. 2016. [A large scale corpus of Gulf Arabic](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 4282–4289, Portorož, Slovenia. European Language Resources Association (ELRA).
- Rania Kora and Ammar Mohammed. 2019. [Corpus on Arabic Egyptian tweets](#).
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. [Similarity of neural network representations revisited](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3519–3529. PMLR.
- Kathrein Abu Kwaik, Motaz Saad, Stergios Chatzikyriakidis, and Simon Dobnik. 2018. [A lexical distance study of arabic dialects](#). *Procedia Computer Science*, 142:2–13. Arabic Computational Linguistics.
- Wuwei Lan, Yang Chen, Wei Xu, and Alan Ritter. 2020. [An empirical study of pre-trained transformers for Arabic information extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4727–4734, Online. Association for Computational Linguistics.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastopoulos, Patrick Littell, and Graham Neubig. 2019. [Choosing transfer languages for cross-lingual learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3125–3135, Florence, Italy. Association for Computational Linguistics.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hanane Nour Moussa and Asmaa Mourhir. 2023. [Darnercorp: An annotated named entity recognition dataset in the moroccan dialect](#). *Data in Brief*, 48:109234.
- Hala Mulki, Hatem Haddad, Chedi Bechikh Ali, and Halima Alshabani. 2019. [L-HSAB: A Levantine Twitter dataset for hate speech and abusive language](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 111–118, Florence, Italy. Association for Computational Linguistics.
- Jonathan Owens. 2006. *A Linguistic History of Arabic*. Oxford University Press.
- Fred Philippy, Siwen Guo, and Shohreh Haddadan. 2023. [Identifying the correlation between language distance and cross-lingual transfer in a multilingual representation space](#). In *Proceedings of the 5th Workshop on Research in Computational Linguistic Typology and Multilingual NLP*, pages 22–29, Dubrovnik, Croatia. Association for Computational Linguistics.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. [Information-theoretic probing for linguistic structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4609–4622, Online. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Faisal Qarah. 2024a. [Egybert: A large language model pretrained on egyptian dialect corpora](#). *Preprint*, arXiv:2408.03524.
- Faisal Qarah. 2024b. [Saudibert: A large language model pretrained on saudi dialect corpora](#). *Preprint*, arXiv:2405.06239.
- Djamé Seddah, Farah Essaidi, Amal Fethi, Matthieu Futral, Benjamin Muller, Pedro Javier Ortiz Suárez, Benoît Sagot, and Abhishek Srivastava. 2020. [Building a user-generated content North-African Arabizi treebank: Tackling hell](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1139–1150, Online. Association for Computational Linguistics.
- Neha Sengupta, Sunil Kumar Sahu, Bokang Jia, Satheesh Katipomu, Haonan Li, Fajri Koto, William Marshall, Gurpreet Gosal, Cynthia Liu, Zhiming Chen, Osama Mohammed Afzal, Samta Kamboj, Onkar Pandit, Rahul Pal, Lalit Pradhan, Zain Muhammad Mujahid, Massa Baali, Xudong Han, Sondos Mahmoud Bsharat, and 13 others. 2023. [Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models](#). *Preprint*, arXiv:2308.16149.
- Skyler Seto, Maartje Ter Hoeve, Richard He Bai, Natalie Schluter, and David Grangier. 2025. [Training bilingual LMs with data constraints in the targeted language](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 19096–19122, Vienna, Austria. Association for Computational Linguistics.
- Guokan Shang, Hadi Abdine, Ahmad Chamma, Amr Mohamed, Mohamed Anwar, Abdelaziz Bounhar, Omar El Herraoui, Preslav Nakov, Michalis Vazirgiannis, and Eric P. Xing. 2025. [Nile-chat: Egyptian](#)

language models for Arabic and Latin scripts. In *Proceedings of The Third Arabic Natural Language Processing Conference*, pages 306–322, Suzhou, China. Association for Computational Linguistics.

Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.

Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R. Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.

Kees Versteegh. 2014. *The Arabic Language*, ned - new edition, 2 edition. Edinburgh University Press.

Zirui Wang, Zachary C. Lipton, and Yulia Tsvetkov. 2020. [On negative interference in multilingual models: Findings and a meta-learning treatment](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4438–4450, Online. Association for Computational Linguistics.

Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 833–844, Hong Kong, China. Association for Computational Linguistics.

Shijie Wu and Mark Dredze. 2020. [Are all languages created equal in multilingual BERT?](#) In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 120–130, Online. Association for Computational Linguistics.

A Probing Datasets

Table 2 presents the datasets we use for probing, where we specify the dialects they support and the dialects we used them for.

Task	Dataset	Dataset Dialect	Used for Dialect	Total Samples	Dialect Samples
POS	NArabizi (Seddah et al., 2020)	Algeria	Algeria	1279	731
	QCRI Arabic POS Dialect (Darwish et al., 2018)	Egypt, Levant, Gulf, Maghreb	Egypt, Morocco	350, 350, 350, 350	338, 145
	Shami (Abu Kwaik et al., 2018)	Levant	Lebanon, Jordan	1069	240, 159
	CAMEL Treebank (Habash et al., 2022)	MSA and CA	MSA	5755 (MSA Variant, 21st Century)	2399
	GUMAR (Khalifa et al., 2016)	Gulf	Saudi, Oman	15205	1774, 1633
NER	CLEANANERcorp (AlDuwais et al., 2024)	MSA	MSA	4899	3050
	ACDNER (Elkhibir et al., 2023)	Egypt, Morocco, Syria	Egypt	353, 378, 361	210
	DzNER (Dahou and Cheragui, 2023)	Algeria	Algeria	5859	1153
	DarNERcorp (Moussa and Mourhir, 2023)	Morocco	Morocco	2511	2177
	Wojood (Jarrar et al., 2022)	MSA, Lebanon, Jordan	Lebanon, Jordan	24643, 383, 4705	190, 570
	MADAR (Tagged) (Bouamor et al., 2018)	Multi-Dialect, MSA	Saudi, Oman	2000 each	2000, 2000
SA	DZYT (Benmounah et al., 2023)	Algeria	Algeria	49942	7828
	AET (Kora and Mohammed, 2019)	Egypt	Egypt	39993	24546
	JHSC (Ahmad et al., 2024)	Jordan	Jordan	276003	30603
	L-HSAB (Mulki et al., 2019)	Lebanon	Lebanon	5377	1142
	LABR (Aly and Atiya, 2013)	MSA	MSA	63257	21696
	MYC (Jbel et al., 2024)	Morocco	Morocco	19455	6306
	MARSA (Alowisheq et al., 2021)	Gulf	Saudi, Oman	68128	10146, 2906

Table 2: Overview of Probing Datasets