

Article

Not peer-reviewed version

ArabianGPT: Native Arabic GPT-based Large Language

[Anis Koubaa](#)*, [Adel Ammar](#), [Lahouari Ghouti](#), [Omer Nekar](#), [Serry Sibae](#)

Posted Date: 26 February 2024

doi: 10.20944/preprints202402.1409.v1

Keywords: Large Language Models, Natural Language Processing, Transformers, Arabic Language, Deep Learning



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

ARABIANGPT: NATIVE ARABIC GPT-BASED LARGE LANGUAGE MODEL

Anis Koubaa, Adel Ammar, Lahouari Ghouti, Omar Najar and Serry Sibae

Robotics and Internet-of-Things Lab, Prince Sultan University, Riyadh

* Correspondence: {akoubaa, aammar, lghouti, onajar, ssibae}@psu.edu.sa

Abstract: The predominance of English and Latin-based large language models (LLMs) has led to a notable deficit in native Arabic LLMs. This discrepancy is accentuated by the prevalent inclusion of English tokens in existing Arabic models, detracting from their efficacy in processing native Arabic's intricate morphology and syntax. Consequently, there is a theoretical and practical imperative for developing LLMs predominantly focused on Arabic linguistic elements. To address this gap, this paper proposes ArabianGPT, a series of transformer-based models within the ArabianLLM suite designed explicitly for Arabic. These models, including ArabianGPT-0.1B and ArabianGPT-0.3B, vary in size and complexity, aligning with the nuanced linguistic characteristics of Arabic. The AraNizer tokenizer, integral to these models, addresses the unique morphological aspects of Arabic script, ensuring more accurate text processing. Empirical results from fine-tuning the models on tasks like sentiment analysis and summarization demonstrate significant improvements. For sentiment analysis, the fine-tuned ArabianGPT-0.1B model achieved a remarkable accuracy of 95%, a substantial increase from the base model's 56%. Similarly, in summarization tasks, fine-tuned models showed enhanced F1 scores, indicating improved precision and recall in generating concise summaries. Comparative analysis of fine-tuned ArabianGPT models against their base versions across various benchmarks reveals nuanced differences in performance, with fine-tuning positively impacting specific tasks like question answering and summarization. These findings underscore the efficacy of fine-tuning in aligning ArabianGPT models more closely with specific NLP tasks, highlighting the potential of tailored transformer architectures in advancing Arabic NLP.

Keywords: large language models; natural language processing; transformers; arabic language; deep learning

1. Introduction

A notable theoretical and empirical challenge in natural language processing (NLP) is the underrepresentation of non-English languages in developing large language models (LLMs). This issue is particularly acute for languages with intricate linguistic structures, such as Arabic. The complexity of Arabic, with its rich morphology and diverse dialects, demands specialized attention in LLM development, a necessity often overlooked in the predominantly English-centric NLP landscape. This disparity limits the scope of NLP applications and hinders progress in understanding and processing multilingual contexts. While GPT-1 and GPT-2 [1,2] have laid the groundwork to demonstrate the capabilities of generative pre-trained transformers, their primary focus has been English, creating a gap in the Arabic NLP landscape. This gap becomes more evident when considering the success of models like Hulmona [3] and ARABERT [4], which, despite their advances in Arabic NLP, still grapple with the challenges of adapting models trained predominantly on English data.

Addressing this gap, ArabianGPT, within the broader ArabianLLM initiative, emerges as a targeted response. Unlike conventional LLMs, which are primarily designed to focus on English and exhibit a propensity to include a significant proportion of English tokens even in their multilingual models, ArabianGPT-Base is explicitly developed for the Arabic language. This model diverges from the norm by prioritizing Arabic's linguistic peculiarities in its design and implementation.

The methodology underpinning ArabianGPT-Base involves three critical components: an architectural adaptation, a rigorous training regimen, and a novel tokenization strategy. Firstly, ArabianGPT-Base adapts the proven GPT-2 architecture, tailoring it to accommodate Arabic's syntactic and morphological idiosyncrasies. This adaptation ensures that the model's structure is inherently aligned with the linguistic realities of Arabic, unlike conventional models that retrofit a predominantly English-centric architecture.

Secondly, the training of ArabianGPT-Base is conducted on a carefully curated corpus, specifically chosen to cover a wide range of Arabic text, thus enabling the model to learn from and adapt to the language's complexity and diversity. This aspect of the methodology is particularly critical, as it ensures that the model is exposed to and learns from a representative sample of the language rather than being limited to a subset that might bias its learning.

Lastly, ArabianGPT-Base introduces AraNizer, an innovative tokenizer developed to address the unique challenges of Arabic text segmentation specifically. This tokenizer represents a significant departure from standard tokenization methods, often inadequate for Arabic due to its complex script and morphological richness. AraNizer's design allows for more precise and contextually aware segmentation, a critical factor in enhancing the model's performance on downstream NLP tasks.

By integrating these methodological innovations, ArabianGPT-Base addresses the gap in Arabic NLP and sets a new standard for developing language-specific models. This approach marks a significant shift in the field, moving away from the one-size-fits-all strategy of previous LLMs towards a more nuanced and inclusive future for NLP.

2. Related Works

Trends

The evolution of Generative Pre-trained Transformers (GPTs) represents a significant milestone in Natural Language Processing (NLP), marking a departure from traditional models towards more sophisticated, contextually aware systems. Initiated by GPT-1 [1], the series showcased the efficacy of Causal Language Modeling (CLM) in understanding and generating human-like text. CLM, by predicting subsequent tokens based on previous context, enables deep learning models to capture nuanced linguistic patterns, facilitating both comprehension and production of complex language structures.

Progressing to GPT-3 [5], the field witnessed an exponential increase in model size and the diversification of training corpora. GPT-3, with its 175 billion parameters, leveraged an extensive dataset comprising vast swathes of internet text, setting new benchmarks across a multitude of NLP tasks. This leap in scale introduced advancements in few-shot learning, empowering models to perform tasks with minimal prior examples, a testament to their evolving contextual understanding.

Simultaneously, the advent of transfer learning emerged as a critical strategy for broadening the applicability of GPT models beyond English, addressing the challenges of linguistic diversity and dataset scarcity in languages like Arabic. The complexity of Arabic, with its rich morphological structure and dialectal variations, underscores the need for specialized approaches in data curation and model tuning.

This trajectory of GPT development not only underscores the importance of scale and data diversity in achieving linguistic comprehension but also highlights the ongoing challenges and opportunities in adapting these models for multilingual NLP, particularly for languages with unique syntactic and morphological features.

Existing Arabic LLMs

Bilingual Focus: Models such as Jais & Jais-Chat [6] and AceGPT [7] have explored bilingual training regimes to enhance Arabic NLP capabilities. These approaches, while innovative, often

struggle to fully capture the nuanced morphological and syntactic features of Arabic due to the inherent complexities of balancing bilingual corpora.

Dialectal Emphasis: Efforts by GigaBERT, ARBERT, and MARBERT [8,9] have been directed towards accommodating the rich dialectal diversity within the Arabic language. Despite their successes, these models may not fully address the intricacies of Modern Standard Arabic (MSA) and classical Arabic, limiting their applicability across all Arabic linguistic tasks.

Understanding/Extraction Focus: ARBERT, AraELECTRA, and CAMELBERT [4,10,11] have set new benchmarks in Arabic NLP for tasks centered around understanding and information extraction. However, their generative capabilities remain limited, underscoring a gap that ArabianGPT seeks to fill.

Text-to-Text Focus: Models like AraT5, QARiB, and JABER & SABIR [12–14] have demonstrated proficiency in text-to-text transformations. Nevertheless, their performance in generative tasks requiring deep linguistic comprehension of Arabic remains an area for further exploration.

AraNizer's Impact

The AraNizer tokenizer stands as a testament to ArabianGPT's commitment to addressing the nuanced challenges of Arabic morphology. By intricately understanding and processing the unique morphological features of Arabic, AraNizer directly enhances ArabianGPT's ability to engage in more accurate and linguistically faithful text generation and processing tasks, setting it apart from the limitations noted in previous models.

ArabianGPT's Positioning

ArabianGPT distinguishes itself by meticulously addressing the gaps left by existing Arabic LLMs. Through a deliberate design that emphasizes Arabic fluency and morphological understanding, ArabianGPT offers unparalleled capabilities in generative tasks. Its architecture, informed by the shortcomings of prior models, is specifically engineered to navigate the complexities of Arabic syntax and morphology, thereby facilitating advanced generative applications in native Arabic contexts, a leap forward in the pursuit of linguistic inclusivity and technological advancement in Arabic NLP.

ArabianGPT represents a significant stride towards addressing the multifaceted challenges of Arabic NLP. By harnessing specialized techniques such as the AraNizer tokenizer and focusing on the generative aspects of language processing, ArabianGPT not only bridges existing gaps but also paves the way for future innovations in the field.

3. Datasets

3.1. Dataset D1

Dataset D1 represents a carefully curated compilation of contemporary Arabic newspaper articles, embodying an expansive collection of 5.37 million texts. This voluminous dataset, which aggregates to a substantial 15.5 GB, encompasses an extensive lexicon tallying 1.8 billion tokens, thereby providing a robust and comprehensive linguistic foundation for the ArabianGPT's training regimen. This corpus encompasses a diverse array of news articles, each contributing to a broad spectrum of linguistic and cultural narratives spanning various chronological timelines. This diversity is paramount in ensuring that the model is exposed to a wide range of dialectical nuances, idiomatic expressions, and syntactic variations, all of which are intrinsic to modern Arabic vernacular and discourse.

The D1 Corpus is not simply a collection of texts; it is a systematic aggregation aimed at capturing the multifaceted nature of Arabic language and culture as represented in contemporary media. The corpus's vastness and variety provide a rich, contextual environment conducive to training sophisticated models capable of understanding and generating Arabic text with a high degree of linguistic fidelity. By covering a broad temporal range, from the early 2000s to 2014, the dataset

facilitates ArabianGPT's comprehension of temporal linguistic shifts and trends, thereby enhancing its capability to produce pertinent and context-aware text.

This corpus has been strategically segmented and organized to facilitate efficient and effective training processes. Each article within the corpus has been pre-processed and structured to ensure optimal compatibility with the ArabianGPT's learning algorithms. The segmentation of the corpus into discrete articles allows for a granular approach to model training, where the nuances of each piece can be individually learned and integrated into the broader linguistic model.

The D1 Corpus, and consequently the ArabianGPT model, is limited by its knowledge cutoff in 2014. This restricts the model's understanding to the language, culture, and world events only up to that year, omitting subsequent developments and modern lexicon. This temporal limitation can impact the model's relevance and effectiveness in contemporary contexts, as it may produce outputs that are outdated or misaligned with recent trends and events. Users and developers must be aware of this constraint, especially for tasks requiring current knowledge.

3.2. Dataset D2: Composite Linguistic Resource

Dataset D2, employed for the training of ArabianGPT-0.3B, amalgamates three distinct yet complementary data sources. Collectively, these datasets constitute a comprehensive linguistic corpus with a substantial size of 23 GB and an impressive lexicon totaling approximately 3.34 billion (3,348,610,363) tokens. The strategic integration of these datasets is predicated on the premise of fostering a multifaceted and robust linguistic environment that is imperative for training a versatile and nuanced Arabic language model.

Collectively, these datasets create a rich tapestry of linguistic scenarios, spanning from formal to colloquial and from historical to contemporary, thus equipping the ArabianGPT-0.3B model with the versatility needed to understand and generate Arabic text across a myriad of contexts and styles. It is this comprehensive and diverse linguistic environment that lays the groundwork for a robust and adaptable Arabic language model, capable of addressing a broad spectrum of language processing tasks with high efficacy.

4. Tokenization

4.1. Overview

Effective natural language processing (NLP) hinges on accurate and efficient text tokenization, particularly for morphologically rich languages like Arabic. Standard tokenization methods often falter when confronted with the unique script, complex morphology, and extensive diacritics inherent in Arabic. This can lead to misinterpretations and hinder the performance of large language models (LLMs) trained on Arabic data.

To address these challenges, we introduced AraNizer, a custom-trained SentencePiece tokenizer specifically designed for Arabic NLP tasks. Leveraging the strengths of the SentencePiece model and incorporating domain-specific knowledge of Arabic linguistics, AraNizer delivers precise and semantically meaningful tokenization, paving the way for improved LLM performance.

Traditional tokenization techniques for Arabic often struggle with:

- **Word root identification:** Separating prefixes, suffixes, and inflections from the core root morpheme poses challenges for standard tokenizers.
- **Diacritic handling:** Diacritics play a crucial role in Arabic morphology and semantics, yet their inconsistent representation and omission in some datasets can lead to ambiguities.
- **Ligature recognition:** The presence of ligatures, where multiple individual characters combine into a single glyph, further complicates word segmentation.

These limitations can adversely affect LLM performance on Arabic tasks, leading to inaccuracies in sentiment analysis, text classification, and other downstream applications. AraNizer aims to overcome these hurdles by:

- **Fine-grained tokenization:** Utilizing the subword-level capabilities of SentencePiece, AraNizer achieves a deeper granularity than simple word segmentation, preserving morphological and semantic information within tokens.
- **Diacritic-aware processing:** AraNizer incorporates diacritics into its tokenization process, ensuring their contribution to word meaning is preserved and reflected in LLM representations.
- **Ligature handling:** AraNizer is trained to recognize and appropriately segment ligatures, avoiding misinterpretations and maintaining accurate morpheme boundaries.

4.2. Description

AraNizer [15] is a custom-trained SentencePiece tokenizer [16] specifically developed for the ArabianLMM series. It utilizes the SentencePiece model, empowering ArabianLMM to parse and generate coherent, contextually relevant Arabic text. The Aranizer tokenizer addresses the complexities of Arabic text with remarkable effectiveness. It leverages SentencePiece and Byte Pair Encoding (BPE) [17] methodologies to optimize tokenization accuracy and efficiency specifically for the Arabic language. This tailored approach ensures optimal handling of the language's morphological richness and unique script challenges. To address the varying complexities and nuances of Arabic text, Aranizer is available in four distinct versions with vocabulary sizes of 32K, 50K, 64K, and 86K. This flexibility allows users to select the version that best suits their specific needs and application domains. For training ArabianGPT models, we opted for the 64K version, striking a balance between accuracy and efficiency.

5. Small-scale model: ArabianGPT-0.1B

5.1. Model Overview

The key objective of the unsupervised pre-training process is to train a language model that excels at predicting upcoming words in a text sequence, without requiring any labeled training data. The dataset D1, described in section 3.1, is used as the unsupervised corpus. ArabianGPT model architecture employs a multi-layer transformer decoder, a powerful neural network architecture adept at capturing long-range dependencies within text. The Transformer architecture has proven highly effective in various language modeling tasks due to its ability to capture long-range dependencies. The transformer decoder applies multi-headed self-attention to focus on relevant relationships between words within the context. The embedding layer transforms the input tokens $U =$ into numerical representations using token and position embeddings. The objective function consists in maximizing the likelihood of predicting the next word in each sequence, given its preceding words. The objective function is expressed as:

$$L(U) = \sum_i \log P(U_i | U_{i-k}, \dots, U_{i-1}; \theta) \quad (1)$$

Where k is the context window size, and θ represents the model parameters (neural network's weights) that are iteratively updated using stochastic gradient descent. In the ArabianGPT model architecture, the employed loss function characterizes the maximization of the conditional likelihood for sequential word prediction, thereby facilitating the model's proficiency in synthesizing contextually coherent and linguistically accurate Arabic text, a quintessential feature for advanced language understanding and generation. The trained model will ultimately generate a probability distribution over potential target tokens using a softmax function. The model's ability to predict upcoming words stems from its learned understanding of language patterns and relationships between words.

Pre-training on a massive, diverse corpus is crucial for the model to capture a broad range of linguistic phenomena.

5.2. Description

ArabianGPT-0.1 follows the architecture of GPT-2-small [2], which is basically equivalent to GPT1 [1] (Figure 1). We retain the underlying principles of the GPT model while incorporating modifications (especially the tokenizer) to better suit Arabic text processing. The ArabianGPT-0.1 model represents a significant advancement in the field of Natural Language Processing (NLP) for Arabic language. Leveraging the GPT-2 architecture, ArabianGPT is custom-tailored to grasp and generate pure Arabic text. Figure 1 shows the main components of this GPT model. It consists of an embedding layer followed by 12 identical layers, each made up of three sub-layers: masked multi-head self-attention, a feed-forward network, and layer normalization. This stacked architecture allows the model to learn complex relationships between words in a sentence:

- **Text and Position Embedding Layer:** Adds information about the individual words and their positions in the sentence, which helps the model to better understand the meaning of the text.
- **Masked Multi-Head Self-Attention:** The core component of the GPT-1 architecture. It allows the model to attend to different parts of the input sentence simultaneously, using multiple heads that focus on different aspects of the words. This gives the model the ability to capture long-range dependencies and context in the text.
- **Feed-Forward Sub-Layer:** Adds non-linearity to the model, allowing it to learn more complex relationships between words.
- **Normalization Layer:** Helps to stabilize the training process and improve the model’s overall performance.




Figure 1. Architecture of ArabianGPT-0.1 model based on GPT2-small.

The ArabianGPT model inherits several key features from the GPT-2 model:

- **Unidirectional Language Modeling:** ArabianGPT, like GPT-2, predicts the probability of a sequence of words in a text, learning to generate coherent and contextually relevant text based on the input it receives.
- **Transformer-based Architecture [18]:** The use of 12 Transformer layers aids in capturing long-range dependencies in text, a crucial aspect for Arabic text which often involves complex sentence structures.
- **Attention Mechanism:** Each Model Attention Layer in ArabianGPT is designed to focus on different parts of the input text, allowing the model to generate more nuanced and accurate Arabic text.
- **Adaptive Learning Rate:** ArabianGPT leverages an adaptive learning rate during training, which facilitates efficient and effective learning from the large-scale Arabic corpus.

5.3. Specifications

Table 1 shows the main characteristics and hyperparameters of the ArabianGPT-0.1 model.

Table 1. Characteristics and hyperparameters of the ArabianGPT-0.1 model.

Attribute	Value
Model Name	ArabianGPT-0.1
Architecture	GPT-2 (GPT2LMHeadModel)
Number of Layers	12
Number of Model Attention Layers (MALs)	12
Number of Heads	12
Vocabulary Size	64,000
Tokenizer	Aranizer
Model Size	134M parameters
Initial Learning Rate	4×10^{-5}
Final Learning Rate	$4 \times 10^{-0.4}$
Context Window Size	768 tokens
Dropout Ratio for Attention (attn_pdrop)	0.1
Dropout Ratio for Embeddings (embd_pdrop)	0.1
Dropout Probability for all Fully Connected Layers (resid_pdrop)	0.1

5.4. Training procedure

ArabianGPT-0.1B was trained on dataset D1 (see section 3.1) using 2 NVIDIA A100 GPUs with 80GB each. The A100 GPUs are known for their efficiency and capability in handling large-scale deep learning tasks. The model underwent training with a dataset comprising 7.5 million sequences, each with a sequence length of 768 tokens. The training parameters included a batch size of 512, a total of 313,500 steps, and a duration of 3 days, resulting in a final loss of 3.97. The training was focused on adapting the model to recognize and interpret the nuances and complexities inherent in Arabic. Table 2 summarizes the conditions and parameters of the training, and Figure 2 shows the training loss convergence through the steps of the training.

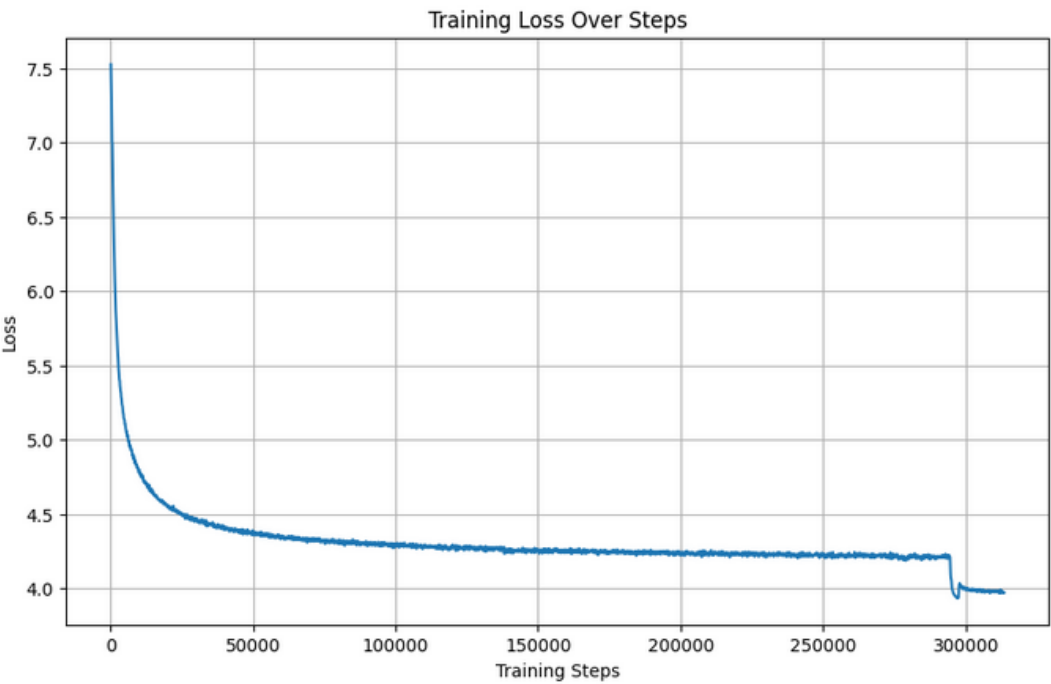


Figure 2. Evolution of the training loss for the ArabianGPT-0.1 model.

Table 2. Details of the training of ArabianGPT-0.1B.

Parameter	Value
Model	ArabianGPT-0.1B (Base Model)
Hardware	2 NVIDIA A100 GPU (80GB each)
Number of Examples	7.5 million sequences (sequence length = 768 tokens)
Batch Size	512
Number of Steps	313,500
Number of Epochs	6
Training Duration	3 days
Final Loss	3.97

5.5. Testing Examples

Table 8 shows some testing prompt and completion pairs for the ArabianGPT-0.1B model. We notice that the completions seem natural, with a few punctuation issues.

Table 3. Sample output of the ArabianGPT-0.1B model.

Completion	Prompt
أعلنت وزارة الحج في المملكة العربية السعودية ، انها ستطلق قبل نهاية العام الميلادي 2008 م خدمة " " اركود " " الإلكترونية ، و " " بسر " " عبر موقعها الإلكتروني الجديد ، فيما ستبدأ الوزارة في تشغيل خدمة الهاتف الصوتي المجانية عبر الجوال (الجوال) . ويأتي الاعلان انفاذا لتوجيهات وزير الحج الدكتور عبدالواحد بن محمد بن بكر القاضي بايقاف الخدمة عن طريق الجوال اعتبارا من أول إبريل المقبل ، والتي تهدف الى إتاحة المكالمات الصوتية المجانية للحجاج الراغبين في أداء مناسك الحج ، وذلك في سبيل تيسير اداء فريضة الحج بكل يسر وسهولة ودون أي جهد	أعلنت وزاره الحج في المملكة العربية السعودية
يبدو اليوم جميلا ، ساقوم بتحضير طعام العشاء لهذا العام ، مع دعوة الجميع إلى العودة إلى المنزل أو الاستراحة لقضاء بعض الوقت مع الأطفال على الشاطئ. أما خلال أيام قليلة ، فإن الجو البارد سيستمر ما يقارب الساعة. كما أن الجو البارد سيكون لطيفا ، وهذا ما سيجعل الأولاد في حال مزاج بارد للغاية. لكن قبل ذلك هناك إمكانية للتخلص من البرد بتناول بعض الاغذية المتلجة ، او الكشخات والعصائر المعبية والعصائر الساخنة ، أما الوجبات الساخنة ، فهي لذيذة جدا ، حيث يزداد تركيز الطفل في هذه الفترات على الفاكهة والخضار والحبوب الطازجة ، والخضار الطازجة. والتي تحتوي على سكر السكراب ، وزيت النخيل ، والبرنقال ، وزيت السمسم ، والبرغر ، وعصير الليمون ، وبعض المكسرات اللبناية ، وعصير الليمون ، والعنب ، والزبيب ، والبطاطس. كما أن تناول الطعام مع الأصدقاء في فترة المساء او تناول بعض الخضروات مثل : الخس والملفوف ، والسردين ، والشاي ، والحرجير ، أو الشاي ، وكذلك بعض العصائر الطازجة	يبدو اليوم جميلا ، ساقوم بتحضير
إن التقنيات الحديثة والمتطورة ساهمت في زيادة انتشار استخدام الانترنت في نشر الاخبار ، حيث تعد اليوم اكبر سوق لنشر الاخبار في الشرق اوسط وذلك في غضون السنوات القادمة. حيث شهد نشر ما يصل الى 20 % من الاخبار المنتورة على الانترنت زيادة في الوقت الذي زاد استخدام الهاتف المحمول ، كما ان سرعة نقل الاخبار عبر الانترنت قد ساعد على زيادة النمو في هذا المجال. وذلك يرجع بشكل أساسي إلى الجهود المتعددة المبذولة من قبل أجهزة الاعلام الرقمية في كل من قطر والكويت والمملكة العربية السعودية	إن التقنيات الحديثة
عاصمة المملكة العربية السعودية هي في وقتنا الحاضر ملتقى لجميع المسلمين في العالم، وهي تحضن عاصمة الإسلام والمسلمين، في مكة المكرمة والمدينة المنورة وبمناخ من من الله، وتشهد دائما لله الحمد والمنة الكثير من المشاريع التنموية الملائقة التي تشهد فيها المنطقة مزيدا من المشاريع المستقبلية. وهذه المنطقة الجميلة والتي تضم أكبر عدد من الحجاج والمعتمرين، تعتبر أكثر مدينة عربية في آسيا من حيث السكان، كما تتميز بتوافر جميع السلع والخدمات الاجتماعية والتعليمية والترفيهية، مما يساهم في تطور كل منطقة من المناطق، والتي تشكل في مجملها نسبة 90% من مجمل مساحة المملكة.. وهذا العدد غير في الواقع، إلا أن المملكة ليست الوحيدة في الغرب، بل يوجد فيها أكبر عدد من الحجاج والمعتمرين، ولا شك أن هناك مواطنين متميزين هم بين الأشقاء، مما يحفز مفهوم الترابط والتكامل بينهم، وهم يحظون بكافة الحقوق، إضافة إلى تمتعهم بكافة الحقوق في كل بقعة من بقاع الأرض، مما يعد أكبر دليل على وجود الإسلام في هذه القارة العظيمة وما تتمتع به من مكانة رائدة بين دول العالم، ناهيك عن المكانة التي يحتلها خادم الحرمين الشريفين لدى المسلمين في مختلف أرجاء المعمورة.. والمملكة العربية السعودية تمثل نموذجا رائعا في هذا المجال.. وهي أرض الإسلام، التي تحضن الكثير من الأديان، وتقوم على التسامح والاحترام، وهذه نعمة على كل مسلم ومسلمة، بفضل الله، بفضل من الله سبحانه وتعالى	عاصمة المملكة العربية السعودية هي

6. Medium-scale model: ArabianGPT-0.3B

6.1. Model Overview

Compared to its predecessor, ArabianGPT-0.1B, ArabianGPT-0.3B is significantly larger and more advanced, trained on a vastly increased dataset of examples. This allows it to capture the intricacies and nuances of Arabic across various domains and topics with greater accuracy and effectiveness.

6.2. Description

ArabianGPT-0.3B is a variant of the GPT-2 model [2] with a medium size, specifically pre-trained for the Arabic language. It retains the underlying principles of GPT-2 while incorporating modifications to better suit Arabic text processing. ArabianGPT-0.3B is a more advanced and larger version than ArabianGPT-0.1. It handles a significantly larger number of examples, making it more adept at capturing the nuances and complexities of the Arabic language across various topics.

6.3. Specifications

Table 4 shows the main characteristics and hyperparameters of the ArabianGPT-0.3 model. It has a double number of layers compared to ArabianGPT-0.1, with a larger number of Model Attention Layers (15 compared to 12) and heads (15 compared to 12).

Table 4. Characteristics and hyperparameters of the ArabianGPT-0.5 model.

Attribute	Value
Model Name	ArabianGPT-MID
Architecture	GPT-2 (GPT2LMHeadModel)
Number of Layers	24
Number of Model Attention Layers (MALs)	16
Number of Heads	16
Vocabulary Size	64,000
Model Size	345M parameters
Initial Learning Rate	4×10^{-5}
Final Learning Rate	$1.1 \times 10^{-0.4}$
Context Window Size	1024 tokens
Dropout Ratio for Attention (attn_pdrop)	0.1
Dropout Ratio for Embeddings (embd_pdrop)	0.1
Dropout Probability for all Fully Connected Layers (resid_pdrop)	0.1

6.4. Training Procedure

Table 5 summarizes the conditions and parameters of the training of the ArabianGPT-0.3B model. Utilizing the power of 4 NVIDIA A100 GPUs, the model underwent an extensive training process, leveraging a dataset consisting of 70 million sequences with a sequence length of 1024. With a batch size of 512 and 1,009,500 steps over 4.97 epochs, the training duration spanned 4.1 days. The culmination of this training journey resulted in a final loss of 3.82, as illustrated in Figure 3, depicting the convergence of the training loss throughout the steps of the training process.

Table 5. Details of the training of ArabianGPT-0.3B.

Parameter	Value
Model	ArabianGPT-0.3B (Medium Model)
Hardware	4 NVIDIA A100 GPU (80GB each)
Number of Examples	70 million sequences (sequence length = 1024)
Batch Size	512
Number of Steps	1,009,500
Number of Epochs	4.97
Training Duration	4.1 days
Final Loss	3.82

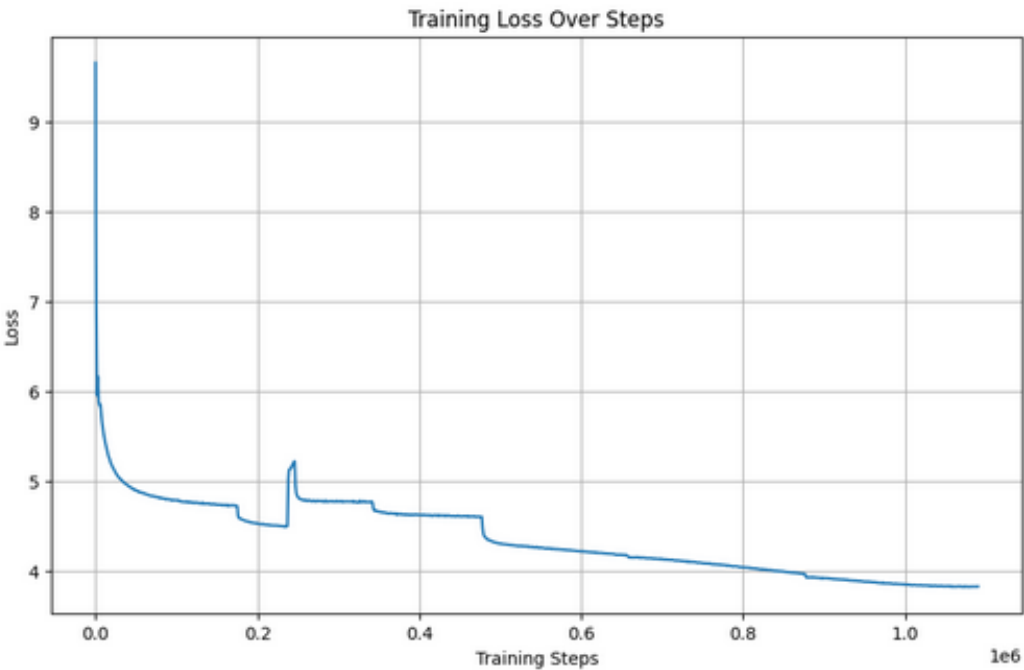


Figure 3. Evolution of the training loss for the ArabianGPT-0.3 model.

6.5. Testing examples

Table 9 shows some testing prompt and completion pairs for the ArabianGPT-0.1B model. We notice that the completions seem natural, with better punctuation.

Table 6. Sample output of the ArabianGPT-0.3B model.

Completion	Prompt
أعلنت وزارة الحج في المملكة العربية السعودية انها تنوي تنفيذ عدة مشاريع كبرى لخدمة حجاج بيت الله الحرام خلال موسم حج هذا العام من ضمنها مشروعات تطوير الفنادق والشقق المفروشة عبر عدة مشاريع ضخمة في كل من مكة المكرمة والمدينة المنورة والمدينة المنورة. هذا المشروع يعتبر من أكبر المشاريع التي أعلنت وزارة الحج عن نتائجها مؤخرا. المشروع يشمل عدة مشروعات فندقية وسياحية والتي تم الانتهاء منها في العام 2017 م. مشاريع أخرى سيتم إعلانها مستقبلا خلال هذا العام وسيتم الإعلان عنها بمشيئة الله تعالى مع بداية موسم حج هذا العام. هذا المشروع هو مشروع سياحي كبير ذو رؤية استراتيجية بعيدة المدى وهو موجه لجميع زوار مدينة بيت الله الحرام وذلك من خلال عدة مشروعات سياحية وسياحية وتجارية وترفيهية وسياحية تستهدف الزوار والمعتمرين القادمين لأداء فريضة الحج من جميع أنحاء العالم	أعلنت وزارة الحج في المملكة العربية السعودية
يبدو اليوم جميلا ، سأقوم بتحضير بعض أطباق الطبخ الشهية لكي تتناسب مع هذا البرج ، وذلك بعد قراءة هذا المقال قبل موعد الزيارة ، والذي سنتحدث فيه عن أفضل برج على الإطلاق. شاهدي هذا المقال واكتشفي ما يناسبك ويجعلك تزورين برج باجي مختار في ماليزيا. برج باجي مختار هو برج جميل محاط بمجموعة من الحقول الخضراء الشاسعة ، ويعد هذا البرج أكثر الأبراج جمالا في ماليزيا ، ويتميز بألوانه الزاهية ورداءه الأبيض الناصع البياض ، بالإضافة إلى ذلك ، فهو واحد من أضخم أبراج العالم. البرج هو برج من صنع الإنسان ، ويتميز برج باجي مختار بلونه الأحمر الفاتح المائل إلى الأبيض ، ويعتبر برج باجي مختار هو أطول أبراج العالم ، حيث يبلغ طوله 17 مترا	يبدو اليوم جميلا ، سأقوم بتحضير
إن التقنيات الحديثة المستخدمة في معالجة الصور الحديثة تستخدم طرقا متطورة لمعالجة الصور بهدف تحسين جودة الصورة ، وهذا ما يفسر سهولة الحصول على الصورة المطلوبة سواء باستخدام برنامج فوتوشوب أو باستخدام برامج الرسم المختلفة التي توفر لك إمكانية التعديل على الصور مثل برنامج " بيكاش	إن التقنيات الحديثة
عاصمة المملكة العربية السعودية هي مدينة الرياض حيث تقع مدينة الرياض على ساحل البحر الأحمر غرب مدينة مكة المكرمة وتقع فوق منطقة وسط السعودية وتحديدا على خط العرض 22 و 24 شمالا وتبعد عن مدينة جدة حوالي 80 كم ، وعلى بعد 130 كم عن مدينة جدة ، وتتميز مدينة الرياض بكونها المدينة الساحلية التي تقع على ساحل البحر الأحمر وتبعد عن العاصمة السعودية الرياض حوالي 38 كم. ويبلغ عدد سكان المدينة حوالي 2.5 مليون نسمة في سنة 2011 م	عاصمة المملكة العربية السعودية هي

7. Few-shot evaluation

In this section, we present a detailed analysis of the performance of ArabianGPT-0.1B, ArabianGPT-0.3B, and compare their performance to AraGPT-Base, AraGPT-Medium, Bloom-7b1, and Llama-7B models across various benchmarks. This evaluation follows a few-shot language model evaluation framework, providing insights into the models' abilities in scientific reasoning, common-sense understanding, cross-domain knowledge, and truthfulness. The benchmarks are designed to test different aspects of language models. They include:

- **ARC (25-shot) normalized:** Assesses scientific reasoning capabilities using normalized metrics.
- **HellaSwag (10-shot):** Evaluates common-sense reasoning in narrative contexts.
- **MMLU (5-shot):** Measures understanding across various subjects.
- **TruthfulQA (0-shot) with mc2:** Gauges the ability to provide truthful and accurate responses.

Table 7 displays the evaluation scores of the aforementioned models on each benchmark. ArabianGPT-0.3B shows a robust aggregate performance across all tasks. It notably excels in TruthfulQA (zero-shot) with a score of 52.5, exceeding Bloom-7b1, and Llama-7B models, which are more than 23 times larger. ArabianGPT-0.1B closely trails its larger counterpart. Despite its smaller size, it competes closely with ArabianGPT-0.3B, particularly in ARC (25-shot) and MMLU (5-shot). The performance efficiency of ArabianGPT-0.1B is notable, given its fewer parameters, indicating effective learning and adaptation capabilities. This demonstrates that smaller models can still deliver competitive performance in diverse linguistic tasks.

Table 7. Zero and Few-shot Evaluation Scores

Model	Language	Average	ARC (25-shot)	HellaSwag (10-shot)	MMLU (5-shot)	TruthfulQA (0-shot)
Bloom-7b1	Multilingual	36.2	31.4	43.3	27.5	42.6
Llama-7B	Multilingual	32.1	24.6	30.9	28.0	45.1
ArabianGPT-0.3B	Arabic	32.7	24.3	28.4	25.7	52.5
ArabianGPT-0.1B	Arabic	31.9	24.0	26.6	25.4	51.8
AraGPT-Base	Arabic	31.7	24.6	27.5	25.1	49.5
AraGPT-Medium	Arabic	32.2	23.9	28.5	26.3	50.0

8. Fine Tuning

8.1. Summarization

The task of summarization in natural language processing (NLP) involves generating concise and meaningful summaries from a larger text. In this project, we have fine-tuned the previously pre-trained ArabianGPT 0.1B and 0.3B Models for the summarization task. These models have been fine-tuned (FT) on a summarization dataset, with the aim of improving their ability to generate accurate and relevant summaries. The performance of these models has been evaluated using the F1 score, a metric that combines precision and recall.

The dataset used for fine-tuning and evaluation is a subset of xlsun [?], specifically xlsun[1000]. This dataset comprises a range of texts suitable for testing summarization capabilities.

The fine-tuning process involves adjusting the weights of the pretrained model specifically for the summarization task. This phase is crucial as it tailors the model to perform better on the desired task, here summarization.

8.1.1. Hyperparameters

In the fine-tuning process of the ArabianGPT models, the Cross-Entropy Loss function was employed alongside the ADAM optimizer. The Cross-Entropy Loss function is widely used in classification tasks, including tasks like summarization and question answering, which can be formulated as predicting the next token in a sequence.

Cross-Entropy Loss: This loss function measures the performance of a classification model whose output is a probability value between 0 and 1. It increases as the predicted probability diverges from the actual label, making it suitable for training models in tasks like summarization and question answering.

Optimizer: ADAM, an algorithm for first-order gradient-based optimization, is utilized for its efficiency in handling sparse gradients and its adaptability to large-scale problems. It’s a popular choice in fine-tuning language models due to its effective handling of large datasets and complex neural network architectures.

Figure 4 depicts the fine-tuning loss for the ArabianGPT-0.1B and ArabianGPT-0.3B models. The smooth decrease in loss values from initial to final stages indicates successful learning and adaptation during fine-tuning. The initial loss values provide a benchmark to gauge the effectiveness of the

fine-tuning process. The significant reduction in these values by the end of fine-tuning reflects the models' improved ability to predict accurate tokens for the tasks at hand. The difference in loss values between the Summarization and Question Answering tasks suggests that these tasks have varying levels of complexity or that the models respond differently to the intricacies of each task.

The ArabianGPT0.3B models start with a higher initial loss in both tasks compared to the 0.1B models, potentially indicating a more complex model capacity which, after fine-tuning, also shows a substantial reduction in loss.




Figure 4. Fine-tuning loss for the ArabianGPT-0.1B and ArabianGPT-0.3B models.

8.1.2. Results

The performance of both the original and fine-tuned models was evaluated using the F1 score. The F1 score is a harmonic mean of precision and recall, providing a balance between the completeness and accuracy of the summaries generated by the models. They are calculated using BLUE and ROUGE Scores using the following equation:

$$F1 = \frac{2 \cdot \text{BLUE} \cdot \text{ROUGE}}{\text{BLUE} + \text{ROUGE}}$$

Figure 5 depicts the F1 score for the ArabianGPT-0.1B and ArabianGPT-0.3B models, before and after fine-tuning on the Summarization dataset. The fine-tuned models (0.1B-FT and 0.3B-FT) demonstrate superior performance compared to their respective non-fine-tuned counterparts. The observed increase in F1 scores post fine-tuning suggests that the models have enhanced their proficiency in summarization tasks. Notably, the 0.3B models, both fine-tuned and non-fine-tuned, exhibit better performance than their 0.1B counterparts, indicating that a larger model capacity can effectively handle the complexity of summarization tasks. Despite these advancements, the F1 scores remain relatively low, signifying substantial room for improvement. Possible avenues for addressing this include further fine-tuning, utilizing a larger and more diverse dataset, or experimenting with different model architectures and hyperparameters.




Figure 5. F1 score for the ArabianGPT-0.1B and ArabianGPT-0.3B models, before and after fine-tuning on the Summarization dataset.

Tables 8 and 9 show a sample output of the fine-tuned ArabianGPT-0.1B and ArabianGPT-0.3B models, respectively, on the summarization task. We notice that the models understand the task and summarize the main idea, but the output sometimes include additional irrelevant words.

Table 10 shows the scores obtained by the fine-tuned ArabianGPT-0.1B and ArabianGPT-0.3B models on the benchmarks introduced in section 7. We notice that the two models have close performance in average. But ArabianGPT-0.3B is significantly better on ARC and HellaSwag benchmarks, while it lags markedly behind ArabianGPT-0.1B on the TruthfulQA benchmark.

Table 8. Sample output of the fine-tuned ArabianGPT-0.1B model on the summarization task.

Completion (ArabianGPT 0.1B)	Prompt
تطورات تكنولوجيا الذكاء الاصطناعي تؤثر على ' الصناعات المختلفة مثل الرعاية الصحية استعراضات".التقدمات جارية".والتصنيع	"لخص النص التالي: يستعرض النص التطورات الأخيرة في مجال الذكاء الاصطناعي وأثرها على الصناعات المختلفة، بما في ذلك الرعاية الصحية والتصنيع. التلخيص:"
التلخيص: استخدام الوسائط المتعددة في تحسين تجربة التعلم للطلاب	"لخص النص التالي: تناول النص أهمية التكنولوجيا في تعزيز التعليم الحديث، مشيراً إلى كيفية استخدام الوسائط المتعددة والأدوات الذكية في تحسين تجربة التعلم لدى الطلاب. التلخيص:"
ن للنص ملخص باللغة العربية: التحديات البيئية " العالمية، التغيرات المناخية والتلوث، وحلول مستدامة للجوانب المعنية	"لخص النص التالي: يتناول النص التحديات التي تواجه البيئة العالمية، مثل التغير المناخي والتلوث، ويقترح حلولاً مستدامة لهذه القضايا. التلخيص:"
التلخيص: دور الإعلام الاجتماعي يؤثر في تشكيل الرأي العام والانتخابات والسياسات العامة. يؤثر أيضا على الانتخابات والاقتصاد. الإعلام الاجتماعي يؤثر	"لخص النص التالي: يناقش النص دور الإعلام الاجتماعي في تشكيل الرأي العام، وكيف يمكن أن يؤثر على الانتخابات والسياسات العامة. التلخيص:"

Table 9. Sample output of the fine-tuned ArabianGPT-0.3B model on the summarization task.

Completion (ArabianGPT 0.3B)	Prompt
<p>لخص النص التالي: يتناول النص التحديات التي تواجه البيئة العالمية، مثل التغير المناخي والتلوث، ويقترح حلولاً مستدامة لهذه القضايا. التلخيص:</p> <p>تلخيص: تلخيص: تلخيص: تلخيص: تلخيص: تلخيص:</p>	<p>لخص النص التالي: يتناول النص التحديات التي تواجه البيئة العالمية، مثل التغير المناخي والتلوث، ويقترح حلولاً مستدامة لهذه القضايا. التلخيص:</p>
<p>التلخيص: يتحدث النص عن التطورات التكنولوجية وأثرها على مختلف القطاعات</p>	<p>لخص النص التالي: يتناول هذا النص بالتفصيل التطورات الحديثة في مجال الذكاء الاصطناعي وتأثيرها على مختلف القطاعات. يبدأ النص بتعريف الذكاء الاصطناعي وأصوله التاريخية، موضحاً كيف تطورت هذه التكنولوجيا عبر العقود. التلخيص:</p>
<p>التلخيص: دراسة الذكاء الاصطناعي: المتغيرات المستمرة والتقدم في التعلم. تلخيص: تلخيص: تلخيص:</p>	<p>لخص النص التالي: يستعرض النص تطور الذكاء الاصطناعي وتأثيره المتزايد على مختلف القطاعات. يبدأ بتاريخ الذكاء الاصطناعي ويتبع تطوره، مع التركيز على التقدم في التعلم الآلي. التلخيص:</p>

Table 10. Model Evaluation Scores for ArabianGPT-0.1B and ArabianGPT-0.3B after fine tuning on the Summarization dataset

Model	Language	Average	ARC (25-shot)	HellaSwag (10-shot)	MMLU (5-shot)	TruthfulQA (0-shot)
ArabianGPT-0.1B-FT	Arabic	31.7	24.5	28.0	25.0	49.3
ArabianGPT-0.3B-FT	Arabic	31.8	27.3	31.5	25.1	43.2

The fine-tuning of ArabianGPT models on the summarization task has demonstrated improvements in their performance as indicated by the F1 scores. The results suggest a positive correlation between model size and summarization ability, as well as the effectiveness of fine-tuning in enhancing model performance for specific tasks. Future work could focus on exploring advanced fine-tuning techniques, larger and more diverse datasets, and possibly integrating other AI methodologies to further enhance the summarization capabilities of transformer-based models.

8.2. Sentiment Analysis

We fine-tuned the Arabic Jordanian General Tweets (AJGT) Corpus, which consists of 1,800 tweets annotated as positive and negative. They are expressed either in Modern Standard Arabic (MSA) or Jordanian dialect. Table 11 details the specifications for the sentiment analysis fine-tuning task.

Table 12 compares the performance of the ArabianGPT-0.1B model on the sentiment analysis task, before and after fine tuning. While the base model gave almost random responses, we notice a marked increase in accuracy after fine tuning attaining 95% after only 3 epochs, on a relatively small fine-tuning dataset.

Table 13 and Figure 6 show a sample output of the ArabianGPT-0.1B model before and after fine tuning on the sentiment analysis task. Before fine tuning, the output scores are close to 0.5, which indicates that the model responses are almost random. By contrast, after fine tuning, the scores are close to 1, and the output is correct in most cases.

Table 11. Training Specifications for the sentiment analysis fine-tuning task.

Parameter	Value
Epochs	3
Learning rate	1.4×10^{-5} to 1.7×10^{-6}
Batch	8
GPU	A100
Global Step	1500
Train Loss	0.203
Training Samples	375
Testing Samples	375
Training Percentage	70%
Testing Percentage	30%
Tokenizer	64K ararizer

Table 12. Model Performance on the sentiment analysis task, before and after fine tuning.

Model	Accuracy
LLM-0.1B-Base	56%
LLM-0.1B-Base (fine-tuned)	95%

Table 13. Sample output of the ArabianGPT-0.1B model before and after fine tuning on the sentiment analysis task.

Fine Tuned Model Results

Index	Prompt	Label	Score	Sentiment
0	المطعم جميل جدا	LABEL_1	0.999522	Positive
1	المطعم سيئ جدا	LABEL_0	0.990597	Negative
2	هذا الكتاب مثير للاهتمام ومفيد	LABEL_1	0.856408	Positive
3	الطعام في هذا المطعم كان فظيحا	LABEL_0	0.982766	Negative

Base Model Results

Index	Prompt	Label	Score	Sentiment
0	المطعم جميل جدا	LABEL_0	0.501228	Negative
1	المطعم سيئ جدا	LABEL_0	0.502999	Negative
2	هذا الكتاب مثير للاهتمام ومفيد	LABEL_1	0.515977	Positive
3	الطعام في هذا المطعم كان فظيحا	LABEL_1	0.570648	Positive



Figure 6. Comparison of the output score of the ArabianGPT-0.1B model before and after fine tuning on the sentiment analysis task, on sample inputs.

8.3. Question Answering

The task of question-answering in natural language processing (NLP) involves generating accurate and contextually relevant answers to given queries. We have fine-tuned the previously pre-trained ArabianGPT 0.1B and 0.3B Models for the question-answering task. These models have undergone fine-tuning (FT) on a question-answering dataset, specifically the mkqa-1000 dataset, with the goal of enhancing their capability to provide precise and informative answers.

The mkqa-1000 dataset, derived from the Google Natural Questions dataset, serves as the basis for fine-tuning and evaluation. This dataset encompasses a diverse set of queries and passage-independent answers, which facilitates testing the models’ question-answering capabilities across various contexts.

The fine-tuning process involves adjusting the weights of the pretrained model, tailoring it specifically for the question-answering task. This crucial phase aims to optimize the model’s performance and enhance its proficiency in generating accurate responses to a wide range of queries in Arabic language.

Table 14 shows the scores obtained by the fine-tuned ArabianGPT-0.1B and ArabianGPT-0.3B models on the benchmarks introduced in section 7. We notice that the two models have close performance in average. But ArabianGPT-0.3B is significantly better on HellaSwag benchmarks, while it lags behind ArabianGPT-0.1B on the TruthfulQA benchmark, suggesting that it was not trained enough with respect to its larger number of parameters.

Table 15 displays a sample output of our two models after fine tuning on the Question/ Answering dataset. While ArabianGPT-0.1B commonly yields irrelevant, incomplete, or unclear responses, some of ArabianGPT-0.3B responses are perfectly accurate, such as for the first and third example.

Table 14. Model Evaluation Scores for ArabianGPT-0.1B and ArabianGPT-0.3B after fine tuning on the Question/ Answering dataset

Model (25-shot) (10-shot) (5-shot) (0-shot)	Language HellaSwag MMLU TruthfulQA	Average	ARC			
ArabianGPT-0.1B-FT	Arabic	31.2	25.2	25.0	24.9	49.6
ArabianGPT-0.3B-FT	Arabic	31.4	24.8	27.6	24.9	48.2

Table 15. Sample output of the two ArabianGPT Models after fine tuning on the Question/ Answering dataset.

Question	ArabianGPT0.3-FT Answer	Comment	ArabianGPT0.1-FT Answer	Comment
ما هي عاصمة مصر؟	عاصمة مصر هي القاهرة.	Accurate answer.	6,449,42720 -202118-202150 -2021 -20215 -2021 -2021 -2021	The response is irrelevant and incorrect.
ما هي عملية البناء الضوئي؟	إذا لم تستطع أيضا أن تتكاثر، هل تستطيع أن تفعل ذلك؟ يجب أن تعرف الخلايا، وخاصة تلك الخلايا التي يوجد بها بعض الجينات التي يمكن أن تكتسب صفات وراثية نتيجة حدوث طفرة ما هي	Response is unclear and incomplete.	يوجد 6 في عدة أقوال ، يمكنك الاستفادة من المواد الكيميائية ، التي تحتوي على نظام ، ومن عليها عملية الزرع. يوجد	The response is unclear, off-topic, and does not answer the question.
ما هو أطول نهر في العالم؟	نهر النيل.	Correct and concise answer.	من 0 حروف ؟ يعتبر من 4 حروف الاجابة '،-': هي	The response is incomplete and does not provide a relevant answer.
ما هي اللغة الرسمية للسعودية؟	يوجد باللغة العربية أكثر من مليون ونصف مليون نسمة، واللغة الإنجليزية هي الثانية في عدد السكان بعد اللغة العبرية، حيث يقطنها نحو 250 مليون نسمة أو أكثر، وهي تحتل المركز	The answer is incorrect and irrelevant.	لا يوجد كلمات متقاطعة المجموعة الثانية من الناحية البشرية، من الناحية الإسلامية.. ولا يوجد فيها أي شيء قدير. في حالة عدم انتشار أي كلمة كلمة أو حساب أو كلمة أو كلمة أو كلمة	The response is irrelevant and does not address the question.

9. Conclusion

The ArabianGPT-Base models represent a significant leap forward in the realm of Arabic natural language processing (NLP), introducing robust capabilities for both understanding and generating Arabic text. The architecture's strength, rooted in its training on a comprehensive corpus and the utilization of the customized tokenizer, AraNizer, collectively contribute to its effectiveness across various NLP tasks specific to the Arabic language context. Notably, the versatility demonstrated by the GPT-2 architecture is further emphasized by the successful adaptation to the intricacies of the Arabic language through the ArabianGPT models.

The development of ArabianGPT is driven by the acknowledgment of a substantial gap in the availability of native Arabic large language models (LLMs), primarily due to the prevalence of English and Latin-based models. Existing models often suffer from the inclusion of English tokens, which

compromises their efficacy in processing the intricate morphology and syntax inherent to native Arabic. This paper introduced ArabianGPT as a tailored solution within the ArabianLLM suite, explicitly designed to address this deficit. The suite includes models of varying sizes, such as ArabianGPT-0.1B and ArabianGPT-0.3B, aligning with the nuanced linguistic characteristics of Arabic.

A key component of the ArabianGPT models is the AraNizer tokenizer, specifically crafted to address the unique morphological aspects of Arabic script. This ensures more accurate text processing, allowing the models to navigate the intricacies of the language with precision. The empirical results from fine-tuning these models on tasks like sentiment analysis and summarization showcase significant improvements. For instance, the fine-tuned ArabianGPT-0.1B model achieved an impressive accuracy of 95%, a substantial increase from the base model's 56%, in sentiment analysis. Similarly, in summarization tasks, fine-tuned models exhibited enhanced F1 scores, indicating improved precision and recall in generating concise summaries.

Comparative analyses of the fine-tuned ArabianGPT models against their base versions across various benchmarks reveal nuanced differences in performance. Fine-tuning is observed to have a positive impact on specific tasks, such as question answering and summarization. These findings underscore the efficacy of fine-tuning in aligning ArabianGPT models more closely with specific NLP tasks, highlighting the potential of tailored transformer architectures in advancing Arabic NLP. As such, the ArabianGPT-Base models not only contribute to bridging the existing gap in Arabic LLMs but also pave the way for further advancements in the field of Arabic NLP.

Funding: This research was funded by Prince Sultan University, grant number SEED-CCIS-2023-145.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets used in this work are publicly available on <https://huggingface.co/datasets>.

Acknowledgments: The authors thank Prince Sultan University for supporting this work.

References

1. Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving language understanding by generative pre-training. *Advances in neural information processing systems*, 2018, Vol. 30.
2. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog* **2019**, 1, 9.
3. ElJundi, O.; Antoun, W.; El Droubi, N.; Hajj, H.; El-Hajj, W.; Shaban, K. hulmona: The universal language model in Arabic. *Proceedings of the fourth Arabic natural language processing workshop*, 2019, pp. 68–77.
4. Antoun, W.; Baly, F.; Hajj, H. Arabert: Transformer-based model for Arabic language understanding. *arXiv preprint arXiv:2003.00104* **2020**.
5. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Amodei, D. Language models are few-shot learners. *Advances in neural information processing systems*, 2020, Vol. 33, pp. 1877–1901.
6. Sengupta, N.; Sahu, S.K.; Jia, B.; Katipomu, S.; Li, H.; Koto, F.; Xing, E. Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models. *arXiv preprint arXiv:2308.16149* **2023**.
7. Huang, H.; Yu, F.; Zhu, J.; Sun, X.; Cheng, H.; Song, D.; Chen, Z.; Alharthi, A.; An, B.; Liu, Z.; others. AceGPT, Localizing Large Language Models in Arabic. *arXiv preprint arXiv:2309.12053* **2023**.
8. Lan, W.; Chen, Y.; Xu, W.; Ritter, A. An empirical study of pre-trained transformers for Arabic information extraction. *arXiv preprint arXiv:2004.14519* **2020**.
9. Abdul-Mageed, M.; Elmadany, A.; Nagoudi, E.M.B. ARBERT & MARBERT: deep bidirectional transformers for Arabic. *arXiv preprint arXiv:2101.01785* **2020**.
10. Antoun, W.; Baly, F.; Hajj, H. AraELECTRA: Pre-training text discriminators for Arabic language understanding. *arXiv preprint arXiv:2012.15516* **2020**.
11. Inoue, G.; Alhafni, B.; Baimukan, N.; Bouamor, H.; Habash, N. The interplay of variant, size, and task type in Arabic pre-trained language models. *arXiv preprint arXiv:2103.06678* **2021**.

12. Nagoudi, E.M.B.; Elmadany, A.; Abdul-Mageed, M. AraT5: Text-to-text transformers for Arabic language generation. *arXiv preprint arXiv:2109.12068* **2021**.
13. Abdelali, A.; Hassan, S.; Mubarak, H.; Darwish, K.; Samih, Y. Pre-training bert on arabic tweets: Practical considerations. *arXiv preprint arXiv:2102.10684* **2021**.
14. Ghaddar, A.; Wu, Y.; Bagga, S.; Rashid, A.; Bibi, K.; Rezagholizadeh, M.; Xing, C.; Wang, Y.; Duan, X.; Wang, Z.; others. Revisiting Pre-trained Language Models and their Evaluation for Arabic Natural Language Processing. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, 2022*, pp. 3135–3151.
15. Najar, O.; Sibae, S.; Ghouti, L.; Koubaa, A. AraNizer 0.1.8. <https://pypi.org/project/aranizer>, 2023.
16. Kudo, T.; Richardson, J. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226* **2018**.
17. Shibata, Y.; Kida, T.; Fukamachi, S.; Takeda, M.; Shinohara, A.; Shinohara, T.; Arikawa, S. Byte Pair encoding: A text compression scheme that accelerates pattern matching **1999**.
18. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems, 2017*, Vol. 30.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.