

BiMediX: Bilingual Medical Mixture of Experts LLM

Anonymous ACL submission

Abstract

In this paper, we introduce BiMediX, the first bilingual medical mixture of experts LLM designed for seamless interaction in both English and Arabic. Our model facilitates a wide range of medical interactions in English and Arabic, including multi-turn chats to inquire about additional details such as patient symptoms and medical history, multiple-choice question answering, and open-ended question answering. We propose a semi-automated English-to-Arabic translation pipeline with human refinement to ensure high-quality translations. We also introduce a comprehensive evaluation benchmark for Arabic medical LLMs. Furthermore, we introduce BiMed1.3M, an extensive Arabic-English bilingual instruction set that covers 1.3 Million diverse medical interactions, resulting in a total of over 632 million healthcare specialized tokens for instruction tuning. Our BiMed1.3M dataset includes 200k synthesized multi-turn doctor-patient chats, and it maintains a 1:2 Arabic-to-English ratio. Our model outperforms state-of-the-art Med42 and Meditron by average absolute gains of 2.5% and 4.1%, respectively, computed across multiple medical evaluation benchmarks in English, while operating at 8-times faster inference. Moreover, our BiMediX outperforms the generic Arabic-English bilingual LLM, Jais-30B, by average absolute gains of 10% on our Arabic medical benchmark and 15% on bilingual evaluations across multiple datasets. Our trained models, instruction set, and source code will be made publicly available.

1 Introduction

Large Language Models (LLMs) demonstrate excellent ability to understand instructions, generate human-like text, and tackle unseen or complex tasks, rendering them valuable for a broad range of applications. Specifically, LLMs have the potential to impact the medical field by offering virtual



Figure 1: BiMediX is adept at conducting precise medical conversations seamlessly in both English and Arabic. It demonstrates advanced interactive capabilities, such as soliciting further details from the patient and seamlessly switch languages based on patient input while maintaining the continuity of the conversation. The Arabic text translates as follows: **User:** *The pain occurs mostly on both sides of the head, sometimes on only one side. It is a throbbing pain that can last for several hours. Light and noise seem to make it worse.* **BiMediX:** *Thank you for the detailed information. Based on your symptoms, it sounds like you may have a migraine. Migraine headaches are usually characterized by a throbbing pain on one or both sides of the head, and can be triggered or increased in intensity by light and noise. It is important to consider that there are other possible causes for your symptoms, so we may need to do further evaluation. Have you noticed any specific factors that trigger these headaches?*

chat support across diverse medical departments to enhance diagnostic accuracy. However, despite the promise of LLMs in healthcare, the medical domain poses distinct challenges, including the necessity for domain-specific data and expertise to ensure accurate interpretations, as well as concerns

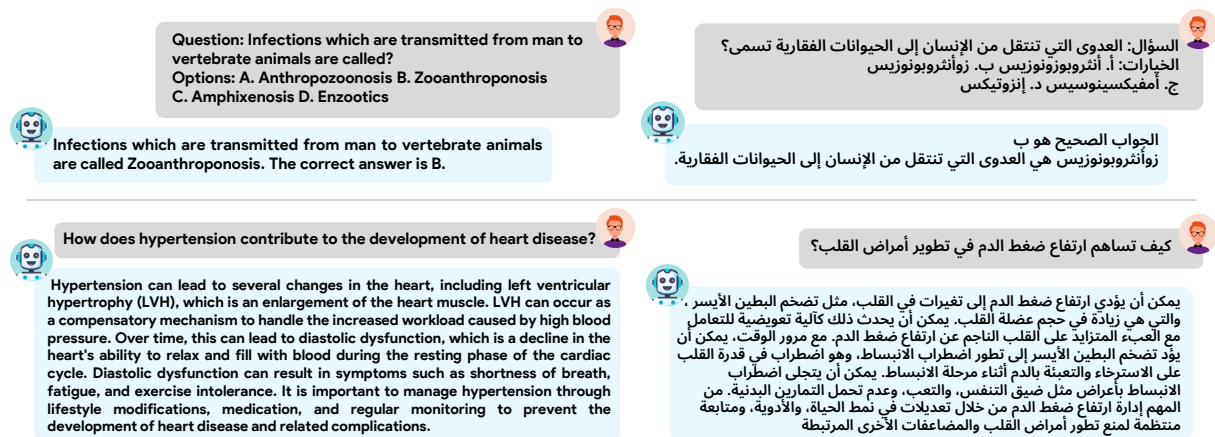


Figure 2: English, Arabic language capabilities of our BiMediX with multiple-choice questions and answers (MCQA) on the top row, and open-ended questions and responses (QA) on the bottom row.

regarding transparency and bias.

Recent methods have emerged to bridge the gap in medical domain knowledge by developing open-source medical LLM alternatives (Chen et al., 2023; Toma et al., 2023; Wu et al., 2023) to proprietary models like ChatGPT (Achiam et al., 2023). However, among the leading medical LLMs in English, Med42-70B (Christophe et al., 2023) has not made its training data and resources available. Meditron-70B (Chen et al., 2023), relies on separate fine-tuning for each evaluation dataset. This approach is likely to compromise the model’s ability to serve as a unified solution with diverse interaction capabilities. Moreover, many of the leading open-source medical LLMs (Christophe et al., 2023; Chen et al., 2023) are limited to the English language, posing limitations in their ability to understand and interact in resource-constrained languages such as Arabic. Despite Arabic’s potential to cater to a population of more than 400 million people, it remains underrepresented in the medical LLM literature. The unique features of Arabic, such as its distinctive script and less conventional right-to-left writing format, the unavailability of large-scale medical training data in Arabic, and the lack of a comprehensive benchmark to evaluate Arabic medical LLMs, present challenges in the development of Arabic or English-Arabic bilingual medical LLMs.

In this work, we introduce BiMediX, a bilingual medical mixture of experts LLM with seamless conversational capabilities in both English and Arabic. Our BiMediX is a chat agent with state-of-the-art performances on non-chat benchmarks in both English and Arabic.

1.1 Contributions

Our key contributions can be summarized as:

- (i) We introduce the *first bilingual medical mixture of experts LLM*, named BiMediX, having seamless interaction capabilities in both English and Arabic languages (see Fig.1). Our model facilitates various medical interactions, including multi-turn chats essential for follow-up inquiries with human patients (Fig.1), multiple-choice question answering (Fig.2 top row) and open-ended question answering (Fig.2 bottom row).
- (ii) We develop a *semi-automated iterative translation pipeline*, incorporating a human verification step to ensure high-quality translation of English medical text into resource-constrained Arabic. This pipeline facilitates the compilation of instruction-tuning dataset and a comprehensive benchmark for evaluating Arabic healthcare LLMs and Arabic-English Bilingual LLMs. The proposed Arabic evaluation benchmark can play a pivotal role in future research endeavors.
- (iii) We curate a comprehensive *Arabic-English bilingual instruction set* named *BiMed1.3M* comprising over 1.3 million instructions, resulting in a total of over 632 million healthcare specialized tokens. Our dataset comprises open-ended question-and-answer, multiple-choice question answering, and over 200k synthesized multi-turn chats rooted in authentic medical content. The integration of multi-turn conversations allows for the creation of a chatbot capable of conducting follow-up inquiries with human patients. Our BiMed1.3M dataset maintains a 1:2 Arabic to English ratio across diverse medical interactions.
- (iv) We introduce a *parameter-efficient fine-tuning*

of routing and expert layers in Mixtral (Jiang et al., 2024) using BiMed1.3M, requiring fewer training resources compared to Med42 and Chen et al. (Christophe et al., 2023; Chen et al., 2023). Our BiMediX achieves *state-of-the-art performance* on multiple medical exam question datasets *in both English and Arabic*. Our model outperforms Med42 and Meditron by average absolute gains of 2.5% and 4.1%, respectively, across multiple medical benchmarks in English, while operating 8 times faster. Additionally, BiMediX surpasses the generic Arabic-English bilingual LLM, Jais-30B, by average absolute gains of 10% on Arabic medical benchmarks and 15% on bilingual evaluations across multiple medical datasets. Moreover, bilingual instruction tuning with our BiMed1.3M leads to an average absolute 10% gain over the base mixtral (Jiang et al., 2024) on bilingual evaluations.

2 Related Works

Proprietary models such as OpenAI’s GPT-3.5 and GPT-4, alongside Google’s Flan-PaLM (Chung et al., 2022), Med-PaLM (Singhal et al., 2022) and Med-PaLM 2 (Singhal et al., 2023) have demonstrated impressive performances in the English medical benchmarks. Nonetheless, the scarcity of public information, alongside concerns regarding access and data privacy, motivate the development of open-source alternatives to proprietary models. These alternatives aim for customization, transparent evaluation, and safe clinical integration.

Early open medical language models, enhanced BERT-style models for medical purposes, leveraging the PubMed corpus for either continuous pre-training (Lee et al., 2020; Huang et al., 2019) or specific domain training from scratch (Beltagy et al., 2019; Gu et al., 2021; Shin et al., 2020). Subsequent research enhanced bidirectional systems to incorporate link structures (Yasunaga et al., 2022b) and knowledge graphs (Yasunaga et al., 2022a). Building on the GPT model series, later studies have tailored GPT-2 based models to medical and scientific literature (Bolton et al.; Luo et al., 2022; Taylor et al., 2022). Recent works such as MedAlpaca (Han et al., 2023) focus on finetuning large scale open-source LLMs such as LLaMA (Touvron et al., 2023a,b) on a set of medical instructions. ChatDoctor (Yunxiang et al., 2023) enhanced a LLaMA model on 100K clinical Q&As, incorporating knowledge retrieval capabilities. Clinical-Camel (Toma et al., 2023) introduced question an-

swering data with the Dialogue-Based Knowledge Encoding (DBKE) process, converting PubMed articles and MedQA into questions and descriptive answers. PMC-LLaMA (Wu et al., 2023) and Meditron (Chen et al., 2023) perform pre-training on PubMed content and medical texts, with further refinements on individual MCQA datasets. Med42 (Christophe et al., 2023) instruction-tuned LLaMA model for medical tasks, though the details of its training remain undisclosed. Additional details on various English medical LLMs provided in Appendix C.

Recent years has witnessed a progress in the Arabic language processing models, including notable models such as AraT5 (Nagoudi et al., 2021) and AraBART (Eddine et al., 2022). In addition to these monolingual models, Arabic has been integrated into multilingual frameworks (Scao et al., 2022; Muennighoff et al., 2022). The recent launch of the Jais model (Sengupta et al., 2023) marked a substantial advancement as a general-purpose bilingual LLM for English and Arabic. Additionally, many of the latest foundational models (Touvron et al., 2023a,b; Jiang et al., 2024) demonstrate the ability to reply in Arabic, despite their limited exposure to the language during pre-training. However, these generic models (Sengupta et al., 2023; Touvron et al., 2023a,b; Jiang et al., 2024) lack sufficient medical domain knowledge, making them ill-equipped for medical applications in both Arabic and English languages.

3 Method

Overview: To develop a bilingual healthcare chat assistant, we first introduce a comprehensive bilingual dataset named BiMed1.3M that encompasses diverse medical interactions such as multiple-choice question answers (MCQAs), open-ended questions, and multi-turn chat conversations, in both English and Arabic (Sec. 3.1). This bilingual dataset is created by first compiling English instruction, as outlined in Sec. 3.1.1. Subsequently, in Sec. 3.1.2 we describe our semi-automated English-to-Arabic translation pipeline, which enables us to obtain high-quality Arabic medical benchmarks and instruction sets, as detailed in Sec. 3.1.3. The generated Arabic instructions are then combined with our English instructions to obtain the BiMed1.3M bilingual dataset. Finally, we present the model and instruction-tuning strategy in Sec. 3.2.

quiries and further information requests. Based on this approach, we produced more than 200,000 high-quality multi-turn medical dialogues, each linked to a specific MCQA, collectively comprising over 74 million tokens.

Following the generation of multi-turn dialogues, our dataset comprehensively encompasses multiple-choice question answering (MCQA), open question answering (QA), and dynamic chat conversations (Chats)—totaling over 860,000 instruction-tuning data in English. Additional details on the diversity and composition of our data are provided in Appendix B. Moreover, we leverage various publicly available medical benchmark datasets in English to evaluate the performance of our model, as elaborated in Sec. 4.2.1.

Next, we describe our semi-automated iterative translation pipeline to generate medical evaluation benchmarks and instruction sets for resource constrained Arabic language.

3.1.2 Semi-Automated Iterative Translation

Our semi-automated iterative translation pipeline is shown in Fig. 4. At first, we perform English to Arabic translation using ChatGPT, which delivers high-quality translations by considering the entire input English text as context. Subsequently, we task ChatGPT with evaluating the quality of the translated Arabic text relative to its original English version. We assign scores ranging from 0 to 100, directly proportional to the translation quality. This scoring aims to quantitatively assess the fidelity and clarity of translations, ensuring the accurate preservation of technical terminology and specifics. Additionally, we verify the alignment of these scores to human preferences with the assistance of native Arabic speakers.

For translations that score below a predefined threshold, we initiate a refinement process. This involves providing ChatGPT with the original English text, along with the current translations and their scores as feedback. We prompt ChatGPT to update the translation to ensure optimal consistency and alignment with the English originals. This feedback loop iteratively improves the translation quality through successive revisions.

However, we observed that certain samples with initially low scores faced limitations in enhancing their scores and translation quality through this automated procedure. To address this, we subject all translations with scores below a threshold to a thorough *manual verification process performed by a*

medical professional fluent in Arabic. To ensure high quality, we additionally sample a random subset of translations with high scores for professional review. Our iterative translation process ensures that translations adhere to rigorous academic and clinical standards.

3.1.3 Bilingual Benchmark & Instruction Set

Creation of Medical Benchmark: To evaluate the accuracy and applicability of Arabic medical AI models, we translated the English medical evaluation benchmarks (detailed in Sec. 4.2.1) into Arabic using the aforementioned iterative translation procedure. By providing a high-quality Arabic medical benchmark aligned with its English counterpart, we aim to establish a fundamental step in bridging the linguistic divide in model evaluation and comparison, thereby offering a valuable asset for further research.

These Arabic benchmarks are combined with the original English evaluation benchmarks to create an English-Arabic bilingual benchmark. This allows us to assess the linguistic capabilities of our bilingual model as well as its medical domain knowledge.

Bilingual Instruction Set: We translated 444,995 English samples into Arabic, covering all three types of medical interactions. Following (Sengupta et al., 2023), we adopted a bilingual approach, mixing Arabic and English in a 1:2 ratio. This approach led to the creation of an extensive bilingual instruction tuning dataset that integrates both languages. Consequently, we compiled an English-Arabic bilingual dataset named BiMed1.3M, 1.5 times larger than our English counterpart, comprising more than 1,311,000 samples. Further details are illustrated in Tab. 2 and Fig. 3.

Next we introduce our model and its bilingual medical instruction tuning.

3.2 Bilingual Medical Instruction Tuning of Mixture of Experts LLM

Recently, Mixtral (Jiang et al., 2024) introduced a Mixture of Experts (MoE) based architecture that achieve impressive performance on generic NLP benchmarks in English. It enables substantial scaling of model size within the same computational budget as traditional dense models. Unlike traditional dense feed-forward network layers, Mixtral employs a set number of "experts" in a sparse manner, replacing the FFN layers in dense networks. Additionally, a gate network or router is employed

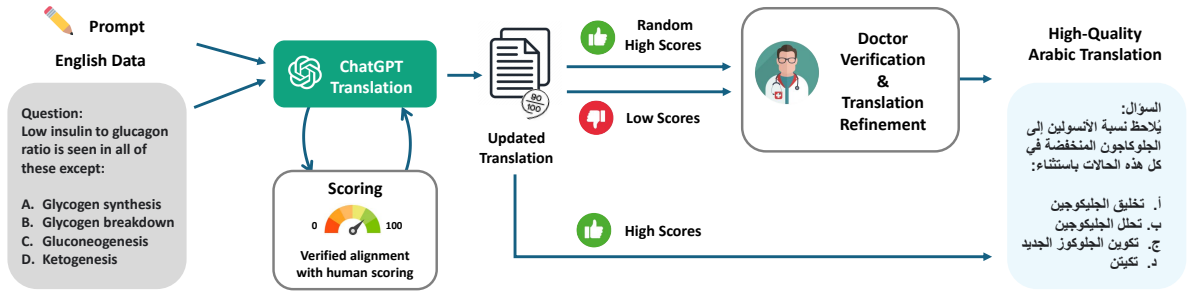


Figure 4: Overview of the proposed semi-automated, iterative translation pipeline featuring human alignment. The process involves iterative translation and score optimization utilizing LLMs (such as ChatGPT) and manual refinement by professionals for samples with low scores along with a random selection of high-scoring samples.

to direct input tokens to appropriate experts based on learned parameters.

While Mixtral (Jiang et al., 2024) offer advantages such as improved inference times over dense models, they face challenges in Arabic-English bilingual medical applications due to limited medical domain knowledge and Arabic language capabilities. To address these challenges, we perform Arabic-English bilingual medical instruction tuning of Mixtral MoE architecture using our BiMed1.3M dataset.

Note that performing conventional supervised fine-tuning of Mixtral (Jiang et al., 2024) requires substantial computational resources and often encounters challenges in carefully tuning the expert layers and the routing network leading to generalization issues. To this end, we employ Parameter-Efficient Finetuning (PEFT) techniques (Hu et al., 2021; Li and Liang, 2021; Dettmers et al., 2023) to adapt the pre-trained Mixtral model (Jiang et al., 2024) with minimal computational resources while ensuring high performance and adaptability. In our approach, we use QLoRA-based PEFT training. Specifically, we attach the QLoRA adapters to the decoder layers of (Jiang et al., 2024), including the experts and the routing network, to repurpose the Mixture of Experts architecture as a bilingual medical chat assistant by utilizing the proposed BiMed1.3M dataset for instruction tuning.

Moreover, to align with the conversational style in our BiMed1.3M dataset, for both English and Arabic instructions, we adopted the Vicuna (Chiang et al., 2023) data format "conversations": ["from": "human/AI," "value": "..."], calculating loss exclusively for the AI's responses within the dialogue script. Unlike alternative approaches such as (Wu et al., 2023), we refrained from conducting

additional medical pre-training before instruction-tuning it on our medical dataset.

In summary, through careful instruction tuning of the routing network and experts in the Mixtral decoder using our BiMed1.3M dataset, we achieve enhanced medical domain capabilities in both Arabic and English languages. This includes proficiency in multiple-choice question answering, open-ended question answering, and realistic multi-turn question answering between our BiMediX model and the user (e.g., patient), making it suitable for deployment as a bilingual medical chatbot.

4 Experiments

4.1 Experimental Settings

Our BiMediX is built upon mixtral (Mixtral-8x7B) (Jiang et al., 2024) base network, a sparse mixture of experts language model. This model extends the Mistral-7B architecture (Jiang et al., 2023), featuring a unique design with each layer hosting eight "expert" feedforward blocks. A router network determines which two experts will process each token, merging their outputs. Consequently, it handles 47 billion parameters, with fewer than 13 billion active during inference. This model features a hidden state dimension of 14,336, a context window capable of accommodating 32,768 tokens, 32 layers, 32 attention heads, and a vocabulary size of 32,000.

For fine-tuning Mixtral, we utilized QLoRA, a low-rank adaptation technique, by incorporating a set of learnable low-rank adapter weights into both the experts and the routing network. This adaptation applies to q , k , and v , setting the rank to 128 and the α parameter to 64. Consequently, this approach results in the training of approximately 4% of the original parameters. Our bilingual model underwent training on roughly 632 million tokens

sourced from the Arabic-English corpus and there are 288 million English tokens to train English model. The training hyper-parameters included an effective batch size of 16 with two gradient accumulation steps and utilized the AdamW optimizer with a learning rate of 0.0002, cosine learning rate schedule, and ten warmup steps. The model implementation leveraged PyTorch (Paszke et al., 2019), along with the Deepspeed and ZeRO (Rasley et al., 2020; Rajbhandari et al., 2020) packages, with gradient checkpointing (Chen et al., 2016). Our BiMedix is trained for two epochs with eight A100 (80GB) GPUs, requiring only 35 hours for the training.

4.2 Quantitive Evaluation

In the literature, evaluating medical language models predominantly involves multiple-choice question-answering tasks, with accuracy as the performance metric. We employed the EleutherAI evaluation framework¹.

4.2.1 Medical Benchmarks

Our analysis incorporates multiple prominent benchmarks in medical multiple-choice question-answering. Two of these datasets are sourced directly from the National Board of Medical Examiners (NBME), responsible for the assessments required for medical licensure in the U.S.

(i) **PubMedQA** (Jin et al., 2019) is a question-answering dataset derived from biomedical research papers on PubMed. Given a question sourced from the title of a research paper and a context from the abstract, the task is to provide an answer in the form of 'yes', 'no', or 'maybe', mirroring the conclusion of the paper. Building on prior research, our examination focuses on the PQA-L subset of PubMedQA, consisting of 1,000 manually annotated QA pairs sourced from PubMed. Answering these questions with accuracy requires thorough reasoning over the biomedical contexts and quantitative data analysis.

(ii) **MedMCQA** (Pal et al., 2022) is a multiple-choice dataset constructed from questions featured in the Indian AIIMS and NEET PG medical entrance exams. It covers a broad spectrum of medical subjects, evaluating both professional domain knowledge and language comprehension. The dataset includes a test set of 4,183 questions, each with four options.

(iii) **MedQA** (Jin et al., 2021) is a dataset comprising multiple-choice questions from medical board examinations in the US, Mainland China, and Taiwan. It features two types of questions: concise sentences to assess specific knowledge and extended paragraphs that detail a patient's condition. Our analysis concentrates on the English portion (USMLE), containing 1,273 samples for testing purposes. Engaging with this benchmark involves multiple stages of reasoning and evidence retrieval.

(iv) **The Medical MMLU** (Massive Multitask Language Understanding) is a group of six datasets that compiles 1,089 test questions with four options related to different medical subjects. The six MMLU datasets are Clinical Knowledge (Cli-KG), College Biology (C-Bio), College Medicine (C-Med), Medical Genetics (Med-Gen), Professional Medicine (Pro-Med) Anatomy (Ana).

All 7045 questions from the above English datasets are translated into Arabic using our semi-automated pipeline to create our Arabic and Bilingual evaluation benchmarks.

4.2.2 Results

Bilingual Evaluation: Here, we evaluate our BiMediX on Arabic-English bilingual evaluation benchmark derived from evaluating results in both languages. Tab. 3 shows our BiMediX results against the base model, Mixtral-8x7B (not fine-tuned) and Jais-30B. BiMediX demonstrates superior performance across all benchmarks, achieving accuracy that is, on average, more than 10 and 15 points higher, respectively, when compared to the baseline model and Jais-30B. This achievement underscores the substantial value of our proposed BiMed1.3M dataset and its unmatched effectiveness and adaptability in addressing medical queries within a bilingual framework.

Arabic Benchmark: We evaluated BiMediX using our Arabic benchmark, comparing its performance with that of Jais-30B, a larger model designed specifically for the Arabic language and capable of functioning in both English and Arabic. In Table 4, we present the findings for Jais and our BiMediX in two configurations: one pre-trained exclusively on Arabic content (Arabic) and the other with bilingual data (Bilingual). Our bilingual model outperforms in all categories within the Arabic context, underscoring that the integration of both types of training data significantly enhances the understanding and processing of medical information in an Arabic setting.

¹<https://github.com/EleutherAI/lm-evaluation-harness>, is the backend for Hugging Face Open LLM Leaderboard.

Model	MMLU						MedMCQA	MedQA	PubmedQA	AVG
	Cli-KG	C-Bio	C-Med	Med-Gen	Pro-Med	Ana				
Jais-30B (Sengupta et al., 2023)	57.4	55.2	46.2	55.0	46.0	48.9	40.2	31.0	75.5	50.6
Mixtral-8x7B (Jiang et al., 2024)	59.1	57.6	52.6	59.5	53.3	54.4	43.2	40.6	74.7	55.0
BiMediX (Bilingual)	70.6	72.2	59.3	74.0	64.2	59.6	55.8	54.0	78.6	65.4

Table 3: Performance of BiMediX on the **Bilingual benchmark**.

Model	MMLU						MedMCQA	MedQA	PubmedQA	AVG
	Cli-KG	C-Bio	C-Med	Med-Gen	Pro-Med	Ana				
Jais-30B (Sengupta et al., 2023)	52.1	50.7	40.5	49.0	39.3	43.0	37.0	28.8	74.6	46.1
Ours (Arabic)	60.0	54.9	55.5	58.0	58.1	49.6	46.0	40.2	76.6	55.4
BiMediX (Bilingual)	63.8	57.6	52.6	64.0	52.9	50.4	49.1	47.3	78.4	56.5

Table 4: Performance of BiMediX on the **Arabic benchmark**.

Model	MMLU						MedMCQA	MedQA	PubmedQA	AVG
	Cli-KG	C-Bio	C-Med	Med-Gen	Pro-Med	Ana				
PMC-LLaMA-13B (Wu et al., 2023)	63.0	59.7	52.6	70.0	64.3	61.5	50.5	47.2	75.6	60.5
Med42-70B (Christophe et al., 2023)	75.9	84.0	69.9	83.0	78.7	64.4	61.9	61.3	77.2	72.9
Clinical Camel-70B (Toma et al., 2023)	69.8	79.2	67.0	69.0	71.3	62.2	47.0	53.4	74.3	65.9
Meditron-70B (Chen et al., 2023)	72.3	82.5	62.8	77.8	77.9	62.7	65.1	60.7	80.0	71.3
BiMediX	78.9	86.1	68.2	85.0	80.5	74.1	62.7	62.8	80.2	75.4

Table 5: Performance of BiMediX on the **English benchmark**.

Model	Active Param	Latency	Tokens/sec
PMC-LLaMA-13B (Wu et al., 2023)	13B	4.1 s	124.1
Med42-70B (Christophe et al., 2023)	70B	24.5 s	20.9
Meditron-70B (Chen et al., 2023)	70B	24.5 s	20.9
Jais-30B (Sengupta et al., 2023)	30B	14.0 s	36.5
Ours-BiMediX	13B	2.8 s	180.6

Table 6: Model statistics comparing memory size and inference speed, computed with (LLM-analysis, 2023) on GPU A100-80GB.

English Benchmark: In this section, we evaluate the quantitative performance of our English model against previous state-of-the-art English medical models. As illustrated in Table 5, BiMediX exhibits outstanding performance across all subsets, securing the highest average scores among all models assessed. When compared to Clinical Camel-70B model, our BiMediX exhibits around 10% average performance gain. Furthermore, it also outperforms Meditron-70B in almost every subset except for MedMCQA. Notably, our model’s success comes without the need for separate fine-tuning on the training set for each evaluation benchmark, in contrast to Meditron, which employs individual fine-tuning for each evaluation benchmark to achieve favorable outcomes. This demonstrates our model’s versatility in handling various medical interactions simultaneously. Moreover, our method outperforms Med42 by an average gain of 2.5%

Furthermore, our model is more efficient in terms of prediction speed (latency and tokens per second) than all other models considered, as detailed in Table 6. Additional qualitative examples are provided in the Appendix D.

5 Conclusion

Our work introduces the first bilingual medical mixture of experts (LLM) BiMediX designed to facilitate comprehensive medical interactions, including conversations, MCQ, and QA in both English and Arabic languages. The implementation of a semi-automated translation pipeline, coupled with human refinement, ensures the high quality of English-to-Arabic translations. BiMediX demonstrates superior performance over existing models such as Med42 and Meditron in English-based medical evaluation benchmarks, and significantly outperforms the generic bilingual LLM, Jais-30B, in Arabic medical and bilingual evaluations. The introduction of a novel evaluation benchmark and the extensive BiMed1.3M instruction set, encompassing 1.3 million diverse medical interactions, further highlights the significance of our contributions.

6 Limitations

Despite its overall improvement, BiMediX, like other language models, may experience issues such

as hallucinations, toxicity, and stereotypes due to inherited limitations from base models and pre-training data. While we’ve conducted automatic and qualitative evaluations with medical professionals, we acknowledge that our model’s medical diagnoses and recommendations may not always be accurate. Extensive human evaluation is more reliable but costlier and time-consuming. The exploration of alternative solutions remains an important focus for ongoing research. Our models lack explicit mechanisms to curb undesirable behaviors. Future research will focus on enhancing security and alignment strategies (Christiano et al., 2017; Bai et al., 2022; Rafailov et al., 2023). On a brighter note, we believe that releasing our weight could contribute to investigating and mitigating risks tied to application uses.

7 Safety and Ethical Implications

We recognize the significant societal impact of this technology, emphasizing ethical considerations and transparency. This release, intended for research, is not ready for clinical or commercial use. Further research is needed to ensure safety and accuracy in clinical settings and to prevent patient harm, in collaboration with patients, medical professionals, and ethicists is crucial for ethical oversight. We adhere to the License Agreement of the base models.

References

Asma Ben Abacha, Eugene Agichtein, Yuval Pinter, and Dina Demner-Fushman. 2017. Overview of the medical question answering task at trec 2017 liveqa. In *TREC*, pages 1–12.

Asma Ben Abacha, Yassine Mrabet, Mark Sharp, Travis R Goodwin, Sonya E Shooshan, and Dina Demner-Fushman. 2019. Bridging the gap between consumers’ medication questions and trusted answers. In *MedInfo*, pages 25–29.

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.

Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. Scibert: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.

Elliot Bolton, David Hall, Michihiro Yasunaga, Tony Lee, Chris Manning, and Percy Liang. BioMedLM: a domain-specific large language model for biomedical text. <https://crfm.stanford.edu/2022/12/15/biomedlm.html>.

Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.

Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, et al. 2023. Meditron-70b: Scaling medical pretraining for large language models. *arXiv preprint arXiv:2311.16079*.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023).

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.

Clément Christophe, Avani Gupta, Nasir Hayat, Praveen Kanithi, Ahmed Al-Mahrooqi, Prateek Munjal, Marco Pimentel, Tathagata Raha, Ronnie Rajan, and Shadab Khan. 2023. Med42 - a clinical large language model.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.

Moussa Kamal Eddine, Nadi Tomeh, Nizar Habash, Joseph Le Roux, and Michalis Vazirgiannis. 2022. Arabart: a pretrained arabic sequence-to-sequence model for abstractive summarization. *arXiv preprint arXiv:2203.10945*.

Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.

Tianyu Han, Lisa C Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K Bresssem. 2023. Medalpaca—an open-source collection of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*.

688	Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan	Ankit Pal, Logesh Kumar Umapathi, and Malaikan-	743
689	Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,	nann Sankarasubbu. 2022. Medmcqa: A large-scale	744
690	and Weizhu Chen. 2021. Lora: Low-rank adap-	multi-subject multi-choice dataset for medical do-	745
691	tation of large language models. <i>arXiv preprint</i>	main question answering. In <i>Conference on Health,</i>	746
692	<i>arXiv:2106.09685</i> .	<i>Inference, and Learning</i> , pages 248–260. PMLR.	747
693	Kexin Huang, Jaan Altosaar, and Rajesh Ranganath.	Adam Paszke, Sam Gross, Francisco Massa, Adam	748
694	2019. Clinicalbert: Modeling clinical notes and	Lerer, James Bradbury, Gregory Chanan, Trevor	749
695	predicting hospital readmission. <i>arXiv preprint</i>	Killeen, Zeming Lin, Natalia Gimelshein, Luca	750
696	<i>arXiv:1904.05342</i> .	Antiga, et al. 2019. Pytorch: An imperative style,	751
697	Albert Q Jiang, Alexandre Sablayrolles, Arthur Men-	high-performance deep learning library. <i>Advances in</i>	752
698	sch, Chris Bamford, Devendra Singh Chaplot, Diego	<i>neural information processing systems</i> , 32.	753
699	de las Casas, Florian Bressand, Gianna Lengyel, Guil-	Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano	754
700	laume Lample, Lucile Saulnier, et al. 2023. Mistral	Ermon, Christopher D Manning, and Chelsea Finn.	755
701	7b. <i>arXiv preprint arXiv:2310.06825</i> .	2023. Direct preference optimization: Your language	756
702	Albert Q Jiang, Alexandre Sablayrolles, Antoine	model is secretly a reward model. <i>arXiv preprint</i>	757
703	Roux, Arthur Mensch, Blanche Savary, Chris Bam-	<i>arXiv:2305.18290</i> .	758
704	ford, Devendra Singh Chaplot, Diego de las Casas,	Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase,	759
705	Emma Bou Hanna, Florian Bressand, et al. 2024.	and Yuxiong He. 2020. Zero: Memory optimizations	760
706	Mixtral of experts. <i>arXiv preprint arXiv:2401.04088</i> .	toward training trillion parameter models. In <i>SC20:</i>	761
707	Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng,	<i>International Conference for High Performance Com-</i>	762
708	Hanyi Fang, and Peter Szolovits. 2021. What disease	<i>puting, Networking, Storage and Analysis</i> , pages 1–	763
709	does this patient have? a large-scale open domain	16. IEEE.	764
710	question answering dataset from medical exams. <i>Ap-</i>	Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and	765
711	<i>plied Sciences</i> , 11(14):6421.	Yuxiong He. 2020. Deepspeed: System optimiza-	766
712	Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W	tions enable training deep learning models with over	767
713	Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset	100 billion parameters. In <i>Proceedings of the 26th</i>	768
714	for biomedical research question answering. <i>arXiv</i>	<i>ACM SIGKDD International Conference on Knowl-</i>	769
715	<i>preprint arXiv:1909.06146</i> .	<i>edge Discovery & Data Mining</i> , pages 3505–3506.	770
716	Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon	Teven Le Scao, Angela Fan, Christopher Akiki, El-	771
717	Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang.	lie Pavlick, Suzana Ilić, Daniel Hesslow, Roman	772
718	2020. Biobert: a pre-trained biomedical language	Castagné, Alexandra Sasha Luccioni, François Yvon,	773
719	representation model for biomedical text mining.	et al. 2022. Bloom: A 176b-parameter open-	774
720	<i>Bioinformatics</i> , 36(4):1234–1240.	access multilingual language model. <i>arXiv preprint</i>	775
721	Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning:	<i>arXiv:2211.05100</i> .	776
722	Optimizing continuous prompts for generation. <i>arXiv</i>	Neha Sengupta, Sunil Kumar Sahu, Bokang Jia,	777
723	<i>preprint arXiv:2101.00190</i> .	Satheesh Katipomu, Haonan Li, Fajri Koto,	778
724	LLm-analysis. 2023. Llm-analysis - latency and mem-	Osama Mohammed Afzal, Samta Kamboj, Onkar	779
725	ory analysis of transformer models for training and in-	Pandit, Rahul Pal, et al. 2023. Jais and jais-chat:	780
726	ference. Availabe at https://github.com/cli99/	Arabic-centric foundation and instruction-tuned open	781
727	llm-analysis .	generative large language models. <i>arXiv preprint</i>	782
728	Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng	<i>arXiv:2308.16149</i> .	783
729	Zhang, Hoifung Poon, and Tie-Yan Liu. 2022.	Hoo-Chang Shin, Yang Zhang, Evelina Bakhturina,	784
730	Biogpt: generative pre-trained transformer for	Raul Puri, Mostofa Patwary, Mohammad Shoeybi,	785
731	biomedical text generation and mining. <i>Briefings</i>	and Raghav Mani. 2020. Biomegatron: Larger	786
732	<i>in Bioinformatics</i> , 23(6):bbac409.	biomedical domain language model. <i>arXiv preprint</i>	787
733	Niklas Muennighoff, Thomas Wang, Lintang Sutawika,	<i>arXiv:2010.06060</i> .	788
734	Adam Roberts, Stella Biderman, Teven Le Scao,	Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mah-	789
735	M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey	davi, Jason Wei, Hyung Won Chung, Nathan Scales,	790
736	Schoelkopf, et al. 2022. Crosslingual generaliza-	Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl,	791
737	tion through multitask finetuning. <i>arXiv preprint</i>	et al. 2022. Large language models encode clinical	792
738	<i>arXiv:2211.01786</i> .	knowledge. <i>arXiv preprint arXiv:2212.13138</i> .	793
739	El Moatez Billah Nagoudi, AbdelRahim Elmadany,	Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres,	794
740	and Muhammad Abdul-Mageed. 2021. Arat5: Text-	Ellery Wulczyn, Le Hou, Kevin Clark, Stephen	795
741	to-text transformers for arabic language generation.	Pfohl, Heather Cole-Lewis, Darlene Neal, et al.	796
742	<i>arXiv preprint arXiv:2109.12068</i> .	2023. Towards expert-level medical question an-	797
		swering with large language models. <i>arXiv preprint</i>	798
		<i>arXiv:2305.09617</i> .	799

Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. Galactica: A large language model for science. *arXiv preprint arXiv:2211.09085*.

Augustin Toma, Patrick R Lawler, Jimmy Ba, Rahul G Krishnan, Barry B Rubin, and Bo Wang. 2023. Clinical camel: An open-source expert-level medical language model with dialogue-based knowledge encoding. *arXiv preprint arXiv:2305.12031*.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Pmc-llama: Further fine-tuning llama on medical papers. *arXiv preprint arXiv:2304.14454*.

Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy S Liang, and Jure Leskovec. 2022a. Deep bidirectional language-knowledge graph pretraining. *Advances in Neural Information Processing Systems*, 35:37309–37323.

Michihiro Yasunaga, Jure Leskovec, and Percy Liang. 2022b. Linkbert: Pretraining language models with document links. *arXiv preprint arXiv:2203.15827*.

Li Yunxiang, Li Zihan, Zhang Kai, Dan Ruilong, and Zhang You. 2023. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. *arXiv preprint arXiv:2303.14070*.

A Appendix

B English Dataset Composition

Data	Samples	Avg. Turns	#Tokens
QA	268.2 K	1.00	51.5 M
MCQA	413.6 K	1.00	163.0 M
Chat	184.8 K	4.71	73.6 M
Total	867.8 K	1.79	288.4 M

Table 7: Statistics of English datasets across QA, MCQA, and Chat (with more than one turn of exchanges) for training and testing, totaling 288.42M tokens.

Data	Samples	Avg. Turns	#Tokens
QA	423.8 K	1.00	131.8 K
MCQA	638.1 K	1.00	342.5 M
Chat	249.7 K	4.72	158.0 M
Total	1311.6 K	1.71	632.3 M

Table 8: Statistics of the BiMed1.3M across QA, MCQA, and Chat (with more than one turn of exchanges) totaling 623M tokens and 1.3M samples. This bilingual dataset, derived from translating approximately 50% of the English dataset into Arabic, is 1.5 times larger due to its 2:1 English-to-Arabic content ratio.

C Related Works

Comparison to Previous Work

Similar to ClinicalCamel, our models are fine-tuned on conversational data, with more than one turn of interactions, to significantly improve understanding and response capabilities in medical contexts. Notably, we are the first to integrate all conceivable interaction types, including Q&A, MCQA, and Chat, into a large-scale instruction tuning dataset. Unlike models that undergo continual pre-training on the base model weights, such as PMC-LLaMA and Meditron, our methodology steers clear of this approach due to the considerable demands for additional data collection, extended training periods, and potential to impair the base model’s capabilities. Furthermore, our models adopt Parameter-Efficient Fine-Tuning (PEFT) (Hu et al., 2021; Li and Liang, 2021; Dettmers et al., 2023) techniques to boost model performance efficiently, circumventing the need for substantial

resources. A summary of close works is provided
in Table 9.

D Additional Examples

863

Model	Language	Continual Pretraining	Instruction Tuning		
			PEFT	Data Type	Data Size
ChatDoctor (Yunxiang et al., 2023)	EN	✗	✗	QA	100 K
MedAlpaca (Han et al., 2023)	EN	✗	✓	QA	160 K
PMC-LLaMA (Wu et al., 2023)	EN	✓	✗	QA + MCQA	514 K
Clinical Camel (Toma et al., 2023)	EN	✗	✓	MCQA + Chat	174 K
Med42 (Christophe et al., 2023)	EN	-	-	-	-
Meditron (Chen et al., 2023)	EN	✓	✗	MCQA	370 K
Ours-BiMediX	AR & EN	✗	✓	QA + MCQA + Chat	1312 K

Table 9: Comprehensive comparison with current open-source medical models, highlighting language, training strategy, and data statistics. PEFT indicates whether models have been adapted for specific tasks with minimal parameter updates. Chat means multi-round QA.

How would you differentiate between Autism Spectrum Disorder and Rett Syndrome in a young child, considering the nuances in developmental milestones and neurobehavioral patterns?

Autism Spectrum Disorder (ASD) and Rett Syndrome are both neurodevelopmental disorders, but they have distinct features. In general, ASD is characterized by impaired social interaction and communication, as well as restricted and repetitive behaviors. Rett Syndrome, on the other hand, is a rare genetic disorder that primarily affects females and is characterized by regression of developmental milestones, loss of purposeful hand skills, and development of repetitive hand movements. While both conditions can present with similar symptoms, the regression of developmental milestones and loss of purposeful hand skills are more specific to Rett Syndrome. Additionally, a genetic test known as the methylation-specific PCR for MECP2 can be used to confirm the diagnosis of Rett Syndrome.

In a case of severe pancreatitis, how do you decide between conservative management and surgical intervention, taking into account the patient's overall condition and comorbidities?

The decision between conservative management and surgical intervention in a case of severe pancreatitis takes into account the patient's overall condition, comorbidities, and potential complications. In cases of mild pancreatitis, conservative management is usually sufficient. However, in cases of severe pancreatitis, surgical intervention may be necessary if there is a concern for complications such as necrotizing pancreatitis, infected pancreatic necrosis, or abscess formation. The patient's overall condition and comorbidities should be considered when making this decision.


How would you manage a case of refractory hyperthyroidism in a patient with Graves' disease who is not a candidate for radioactive iodine therapy?

Refractory hyperthyroidism in a patient with Graves' disease who is not a candidate for radioactive iodine therapy can be managed with pharmacological thyroid suppression. This involves using medications to reduce the production and release of thyroid hormones. One commonly used medication is propylthiouracil, which inhibits the synthesis of thyroid hormones. Another medication is methimazole, which also inhibits the synthesis of thyroid hormones. These medications can be used alone or in combination with other treatments, such as beta-blockers, to manage symptoms of hyperthyroidism and control thyroid hormone levels. It is important to work closely with a healthcare provider to determine the most appropriate treatment plan for the individual patient, as the choice of medication and dosage will depend on factors such as the patient's age, overall health, and other medical conditions they may have.

Is surgery a viable option in this case?

Surgery may be considered as a treatment option for refractory hyperthyroidism in a patient with Graves' disease who is not a candidate for radioactive iodine therapy. However, it is important to note that surgery carries a higher risk of complications than other treatment options, such as pharmacological thyroid suppression. Surgical options include thyroidectomy, which involves removing part or all of the thyroid gland, and thyroid lobectomy, which involves removing only one lobe of the thyroid gland. These procedures may be recommended for patients who have failed other treatment options or who have specific contraindications to other treatments. The decision to undergo surgery should be made in consultation with a healthcare provider, taking into account the patient's individual circumstances and the potential risks and benefits of the procedure.


Figure 5: Qualitative Examples of our BiMediX (English).



لقد تم تشخيص إصابتي بمرض السكري من النوع 2. هل يمكنك توضيح كيف يجب أن أراقب مستويات السكر في الدم، وما هي التغييرات الغذائية التي قد تكون ضرورية؟



يُعدُّ مراقبة مستويات السكر في الدم أمرًا حاسمًا للسيطرة على السكري من النوع 2. يُنصح بمراقبة مستوى السكر في الدم مرة واحدة على الأقل في اليوم. يمكن أن تساعد رصد مستويات السكر في الدم في فهم كيف يؤثر نمط حياتك ونظامك الغذائي على مستويات سكر الدم. تتطلب تغييرات في النظام الغذائي، مثل تقليل استهلاك السكر المضاف والكربوهيدرات ذات الامتصاص السريع وزيادة تناول الألياف، عادةً للتحكم في السكري من النوع 2. فمن المهم أيضًا تناول وجبات متوازنة ومنظمة للمساعدة في تنظيم مستويات السكر في الدم. أوصي بمناقشة خطة مراقبة السكر في الدم وتغييرات في النظام الغذائي مع مقدم الرعاية الصحية الخاص بك.



كنت أعاني من ضيق في التنفس والتعب في بعض الأحيان. هل يمكن أن تكون هذه علامات مبكرة لمرض في القلب أو الرئة؟



يمكن أن يكون ضيق التنفس والتعب أعراضًا مبكرة لعدة حالات تتعلق بالقلب والرئة. يمكن أن يكون ذلك مرتبطًا بأمراض القلب مثل القصور القلبي أو أمراض الشرايين التاجية، أو أمراض الرئة مثل الربو أو الانسداد الرئوي المزمن. من المهم التشاور مع مقدم الرعاية الصحية الخاص بك لتقييم أعراضك وإجراء الفحوصات المناسبة لتحديد السبب الأساسي. سيتمكنون من تقديم تشخيص دقيق وتوصية بخطة علاجية مناسبة.



هل هناك أي شيء يمكنني فعله بنظامي الغذائي للمساعدة في تحسين حالتي؟



يمكن أن تكون التغييرات في النظام الغذائي لها تأثير إيجابي على حالات القلب والرئة. عمومًا، يجب تضمين الأطعمة الغنية بالألياف والفواكه والخضروات والأحماض الدهنية أوميغا-3 في النظام الغذائي. قد تشمل الأطعمة التي يجب تقليل تناولها منتجات الألبان المحتوية على نسبة عالية من الدهون المشبعة، واللحوم المصنعة، والمشروبات الغازية، والأطعمة المصنعة. بالإضافة إلى ذلك، قد يوصي مقدم الرعاية الصحية باتباع نظام غذائي محدد مثل النظام الغذائي المتوسط أو نظام غذائي قليل الصوديوم اعتمادًا على حالتك الخاصة. من المهم استشارة مقدم الرعاية الصحية قبل إجراء أي تغييرات كبيرة في النظام الغذائي لضمان تلبية احتياجات

Figure 6: Qualitative Examples of our BiMediX (Bilingual) which is translated to English as follows:

User: I have been diagnosed with type 2 diabetes. Can you explain how I should monitor my blood sugar levels, and what dietary changes might be necessary? **BiMediX:** Monitoring blood sugar levels is crucial to controlling type 2 diabetes. It is recommended to monitor your blood sugar level at least once a day. Monitoring your blood sugar levels can help you understand how your lifestyle and diet affect your blood sugar levels. It usually requires dietary changes, such as reducing consumption of added sugar and fast-absorbing carbohydrates and increasing fiber intake, to control type 2 diabetes. It is also important to eat regular, balanced meals to help regulate blood sugar levels. I recommend discussing a blood sugar monitoring plan and diet changes with your healthcare provider.

User: I've been experiencing occasional shortness of breath and fatigue. Could these be early signs of a heart or lung condition? **BiMediX:** Shortness of breath and fatigue can be early symptoms of several heart and lung conditions. This can be related to heart disease such as heart failure or coronary artery disease, or lung disease such as asthma or chronic obstructive pulmonary disease. It is important to consult with your healthcare provider to evaluate your symptoms and perform appropriate tests to determine the underlying cause. They will be able to provide an accurate diagnosis and recommend an appropriate treatment plan.

User: Is there anything I can do with my diet to help improve my situation ? **BiMediX:** Dietary changes can have a positive impact on heart and lung conditions. Generally, foods rich in fiber, fruits, vegetables and omega-3 fatty acids should be included in the diet. Foods to reduce your intake may include dairy products high in saturated fat, processed meats, soft drinks, and processed foods. In addition, your healthcare provider may recommend a specific diet such as a moderate-moderate diet or a low-sodium diet depending on your specific condition. It is important to consult your health care provider before making any major dietary changes to ensure your needs are met.