

MedPT: A Massive Medical Question Answering Dataset for Brazilian-Portuguese Speakers

Fernanda B. Färber^{1,2}, Iago A. Brito^{1,2}, Julia S. Dollis^{1,2}, Pedro S. F. B. Ribeiro^{1,2}
Rafael T. Sousa^{1,3}, Arlindo R. Galvão Filho^{1,2}

Advanced Knowledge Center for Immersive Technologies (AKCIT)¹

Federal University of Goiás²

Federal University of Mato Grosso³

{fernandabufon, iagoalves, juliadollis, schindler}@discente.ufg.br

rafaelsousa@ufmt.br, arlindogalvao@ufg.br

Abstract

While large language models (LLMs) show transformative potential in healthcare, their development remains focused on high-resource languages. This creates a critical barrier for other languages, as simple translation fails to capture unique clinical and cultural nuances, such as endemic diseases. To address this, we introduce MedPT, the first large-scale, real-world corpus of patient-doctor interactions for the Brazilian Portuguese medical domain. Comprising 384,095 authentic question-answer pairs and covering over 3,200 distinct health-related conditions, the dataset was refined through a rigorous multi-stage curation protocol that employed a hybrid quantitative-qualitative analysis to filter noise and contextually enrich thousands of ambiguous queries, resulting in a corpus of approximately 57 million tokens. We further utilize LLM-driven annotation to classify queries into seven semantic types to capture user intent. To validate MedPT's utility, we benchmark it in a medical specialty classification task: fine-tuning a 1.7B parameter model achieves an outstanding 94% F1-score on a 20-class setup. Furthermore, our qualitative error analysis shows misclassifications are not random but reflect genuine clinical ambiguities (e.g., between comorbid conditions), proving the dataset's deep semantic richness. We publicly release MedPT on [Hugging Face](#) to support the development of more equitable, accurate, and culturally-aware medical technologies for the Portuguese-speaking world.

Keywords: Medical question answering, healthcare, Brazilian Portuguese

1. Introduction

Large language models (LLMs) are rapidly expanding their reach across diverse healthcare applications, from online patient support and clinical text classification to robust medical question answering (Nazi and Peng, 2024; Liu et al., 2025). However, the performance, safety, and clinical reliability of these systems are fundamentally contingent on the scale and quality of their training data (Ryan Conmy et al., 2023). Consequently, datasets that lack demographic representation, contain inherent biases, or are insufficiently curated to handle complex medical nuances can lead to models that produce factually incorrect or systematically skewed outputs, posing significant risks and undermining clinical trust (Khan et al., 2021; Schwabe et al., 2024). Compounding this challenge, current progress and resource development are overwhelmingly concentrated in high-resource languages, particularly English, a linguistic bias that creates a critical barrier for other major world languages, severely limiting the development, validation, and equitable deployment of these promising technologies to large global populations (Enshaei and Naderkhani, 2024).

This linguistic and data-centric gap is clearly reflected in the literature, where significant research

investment in medical text curation has centered on English (Jin et al., 2019; Pal et al., 2022; Singhal et al., 2022; Preston et al., 2023), Chinese (He et al., 2019; Zhang et al., 2018; Zhu et al., 2023), or both (He et al., 2020). Despite valuable initiatives have begun to extend this scope to other languages such as Arabic, Vietnamese, Persian, and Turkish (Abdelhay et al., 2023; Nguyen et al., 2022; Ghasabi et al., 2025; İncidelen and Aydoğan, 2024), the resource disparity remains particularly acute for Brazilian Portuguese in specialized domains, such as medical question answering.

While translating existing datasets may seem like a viable shortcut, this approach is fundamentally limited because it fails to capture essential local context. For example, the epidemiology of endemic conditions such as Dengue and Chagas disease is entirely absent from these resources. Furthermore, translation erases crucial linguistic and cultural nuances, including region-specific clinical expressions and colloquial patient descriptions of symptoms (Naveen and Trojovský, 2024). This creates a dual problem: locally, it constrains the advancement of clinically reliable language technologies for the population. Globally, it perpetuates a critical void in the medical data landscape, which remains heavily biased towards the health contexts prevalent in high-resource language communities.

Therefore, creating a robust dataset for Brazilian Portuguese is not only essential for its speakers but also represents an opportunity to enrich global health knowledge with data on otherwise underrepresented diseases and clinical realities.

In this work, we introduce MedPT¹, the first public large-scale corpus for the Brazilian Portuguese medical domain, comprising 384,095 authentic question-answer pairs from real patient-doctor interactions along with metadata detailing the disease, the medical specialty, and a label identifying the question type. We provide a comprehensive analysis of the dataset, detailing its linguistic properties, such as the natural asymmetry of patient and doctor communication, and its thematic breadth, which covers over 3,200 medical conditions across 104 medical specialties. To validate its practical utility, we establish strong benchmark on a real-world medical speciality classification task, achieving 94% F1-score on a challenging 20-class setup by fine-tuning a modern 1.7B parameter model on MedPT, confirming the dataset is a high-quality and effective resource for developing and specializing language models for the nuanced Brazilian healthcare context.

This resource provides a foundational dataset for a range of tasks, including medical question answering, patient intent classification, and linguistic studies on doctor-patient communication. More broadly, MedPT captures the real-world nuances of healthcare interactions, enabling the development of more contextualized, reliable, and accessible models for Portuguese-speaking populations.

2. Related Works

In recent years, several datasets have been developed to support research in medical question answering and health-related natural language processing (NLP). These resources have been fundamental for training models capable of reasoning over clinical information, understanding patient intent, and generating reliable medical responses. Natural language processing has been applied across multiple medical fields, including oncology, psychiatry, and primary care (Casella et al., 2024; Wang et al., 2020; Wu et al., 2019). Within these domains, the most common NLP tasks are information extraction (Mulyar et al., 2021; Cury et al., 2021; Alkouz et al., 2019), classification (Magna et al., 2020), named entity recognition (Ji et al., 2019) and question answering (Singhal et al., 2025).

However, the development and fine-tuning of such models require large, well-structured, and high-quality datasets. To address this, several datasets have been proposed to advance research

in medical natural language processing. For example, Johnson et al. (2023) introduced MIMIC-IV, which focuses on describing a vast collection of de-identified electronic health records from intensive care units, widely used for developing clinical prediction models and training NLP systems on real-world clinical documentation. Similarly, BioRED (Luo et al., 2022) aimed at advancing biomedical knowledge extraction by annotating document-level relations between diverse entities like genes, diseases, and chemicals from scientific papers abstracts.

Despite these contributions, most existing datasets consist primarily of clinical or biomedical records and lack open-ended, real-world patient questions paired with verified medical answers. He et al. (2019) proposed cMedQA, which includes 54,000 pairs of physician-patient questions and answers in Chinese, sourced from online health forums, focusing on non-restricted, open-ended questions. Similarly, MedRedQA (Nguyen et al., 2023) includes 51,000 pairs of open-ended English consumer questions and verified expert answers, collected from online medical consultation forums.

Moreover, pioneering efforts have begun to establish benchmark medical datasets for several under-resourced languages, employing a variety of curation strategies. One common approach involves collecting data from trusted online health portals. For instance, ViHealthQA (Nguyen et al., 2022) contains over 10,000 medical question-answering pairs for Vietnamese sourced from credible health websites, while MAQA (Abdelhay et al., 2023) represents the largest Arabic healthcare Q&A dataset, with over 430,000 questions from regional health platform. Other initiatives have utilized different sources; MedTurkQuAD (İncidelen and Aydoğın, 2024), a Turkish medical QA dataset, was derived from Turkish Wikipedia and national medical theses. The effort for Persian by Ghassabi et al. (2025) was even more comprehensive, introducing the MF3QA dataset and leveraging it to create Gaokerena-V1, the first open-source Persian medical language model. Together, these works illustrate a growing cross-lingual interest in expanding medical NLP resources beyond English, particularly for morphologically rich or culturally distinct languages.

Recent efforts to address the Portuguese data shortage include the work of Paiola et al. (2024), who fine-tuned LLMs on an English healthcare corpus translated into Brazilian Portuguese. While this approach enabled preliminary evaluations of medical dialogue assistants, the resulting dataset was not made publicly available. More fundamentally, its reliance on translation means it inherently lacks the authentic, region-specific linguistic nuances and clinical contexts, such as endemic diseases, that only a native, real-world dataset can capture.

¹<https://huggingface.co/datasets/AKCIT/MedPT>

This highlights a persistent and critical need for a publicly accessible, native corpus built from authentic patient-doctor interactions to advance the field.

Alongside translation efforts, existing medical datasets for native Brazilian Portuguese have centered on formal clinical text from hospital records. Prominent examples include BRATECA (Consoli et al., 2022), a large collection of de-identified clinical notes; SemClinBr (Oliveira et al., 2022), a semantically annotated corpus for entity recognition; and the work of (da Costa Lopes, 2019), identifying symptoms in hospital texts. While valuable, these resources are composed of formal medical documentation and do not capture the open-ended, conversational dialogues between patients and physicians. This absence of real-world, patient-initiated questions is a primary barrier to developing and evaluating robust interactive healthcare systems in Portuguese.

A recent initiative targeting this conversational gap in Portuguese is IaraMed (Färber et al., 2025), a women’s healthcare chatbot developed using data scraped from online health platform. While a promising application of real-world medical Q&A, this work is narrowly focused on women’s health and does not cover the full spectrum of patient-doctor inquiries, and the underlying dataset was not publicly released. Thus, a critical need persists for a broad-domain, publicly accessible, native corpus built from authentic patient-doctor interactions to advance the field.

3. MedPT

This section details the multi-stage process applied in the development of MedPT dataset, designed to transform raw, user-generated content into a robust language resource. We first describe the data acquisition from an authentic patient-doctor interaction platform, followed by our extensive data cleaning and curation pipeline, which includes both general preprocessing and nuanced, column-specific refinements. We then detail an LLM-driven annotation step, where we employed a large language model to classify each user query and identify its question type (e.g., treatment, diagnosis). To enhance the dataset’s utility. Finally, to ensure a fair comparison and compatibility with modern models, all metrics were extracted using the sub-word tokenizer adopted by state-of-the-art LLMs (Achiam et al., 2023).

3.1. Data Acquisition

To build the MedPT dataset, we collected real medical questions and answers in Portuguese from

Doctoralia², a widely used online platform where Brazilian patients ask health-related questions and receive responses from verified doctors. Doctoralia was chosen as data source due both its large volume of questions across a wide range of medical specialties and its verification process to ensure that only certified health professionals can respond to users’ inquiries. Additionally, all questions are completely anonymous, which guarantees the privacy and confidentiality of users’ personal information. Furthermore, the data collected from this platform are natural and user-generated, reflecting real-life inquiries from Brazilian users. This aspect makes the dataset particularly valuable for developing healthcare solutions tailored to the Brazilian context, as it captures the country’s linguistic and cultural nuances, healthcare system characteristics, and the prevalence of region-specific diseases.

To collect the questions, answers, and related data, we developed a custom web-scraping pipeline to extract information directly from the website. Each sample includes the medical condition of the question (e.g., hypertension, pregnancy, diabetes), the patient’s textual query, the doctor’s textual response, and metadata about the physician, such as specialty and rating. During this collection phase, we gathered approximately 2 million raw samples.

3.2. Data Cleaning

The initial cleaning phase focused on structural normalization and deduplication. A primary source of data redundancy was identified in the platform’s design, which allowed a single question-answer pair to be associated with multiple medical specialties. This resulted in numerous duplicate entries for the same core text. To address this and consolidate the dataset, we removed the redundant specialty column and then applied a strict deduplication process based on the core question-answer text. This procedure, combined with the removal of other irrelevant columns (e.g., doctor’s ratings), reduced the corpus from approximately 2 million raw samples to 393,143 unique entries. Subsequently, we performed a textual normalization step across all columns, removing HTML tags, URLs, special characters, and extraneous whitespace, which significantly reduced noise as evidenced by a decrease in the average token count per question-answer pair from 33.05 to 26.54 tokens.

3.2.1. Column-Specific Curation

To ensure the quality and suitability of the corpus for question-answering tasks, we implemented a multi-stage, column-specific curation protocol through

²<https://www.doctoralia.com.br>

all features of our dataset (*Question*, *Answer Condition* and *Medical Specialty*). This protocol combined quantitative analysis with qualitative manual inspection. We first analyzed token distributions such as mean, standard deviation (SD) and range to identify statistical outliers, which were then manually reviewed to confirm their nature. Based on this hybrid analysis, we applied targeted procedures, including the removal of malformed or uninformative entries and the contextual enrichment of incomplete questions, as detailed below.

Question. The *Question* column exhibited high linguistic diversity (mean 42.5, SD 31.0), with token counts ranging from 3 to 801. Our qualitative inspection confirmed that the longest entries were valid, complex patient narratives. Conversely, the shortest entries (3–9 tokens) were typically found to be context-dependent (e.g., "What is the treatment?"), lacking sufficient information to be standalone queries. To address this ambiguity without data loss, we implemented a contextual enrichment procedure: the corresponding *Condition* was prepended to each of these 7,178 incomplete queries. This process rendered each question self-contained and preserved a significant portion of the dataset that would have otherwise been discarded.

Answer. The *Answer* column exhibited even greater variability (mean 96.5, SD 75.0), with token counts spanning from 0 to 2,984. To ensure the corpus contained only direct, high-quality responses, a qualitative review of these extremes was performed. This analysis revealed two distinct types of low-utility data: answers with fewer than 10 tokens were found to be uninformative (e.g., "Consult a specialist"), while those exceeding 1,000 tokens, though medically valid, were typically generic informational articles rather than specific replies to the patient's query. To optimize the dataset for direct question-answering tasks, we established a valid token range of 10–1,000. Applying this filter removed 8,359 samples, significantly improving the signal-to-noise ratio and focus of the final corpus.

Condition. For the *Condition* column, we implemented a token-count-based filtering protocol to ensure data quality. An initial analysis programmatically removed all zero-token samples. We then conducted a qualitative review of the distribution's extremes. This inspection revealed that queries containing very few tokens (e.g., "anxiety") were valid medical conditions and were therefore retained. Conversely, entries with the highest token counts were found to be consistently malformed or corrupted data. Based on this manual review, we established a conservative upper threshold of 23 tokens to precisely exclude these erroneous entries

while preserving all valid conditions. This procedure refined the column's consistency by removing 398 samples.

Medical Specialty. The *Medical Specialty* column presented a distinct curation challenge. Statistical analysis showed moderate variance (mean 5.7 tokens, SD 3.4), which our qualitative review attributed to the platform's feature allowing professionals to list multiple specialties. We made the deliberate decision to retain these multi-specialty entries to avoid information loss, as they represent valid, rich data. However, a separate filtering process was necessary to address genuine scraping artifacts. Our analysis identified and removed non-specialty data, such as location names (e.g., "Fortaleza"), numerical values, and empty strings. This targeted procedure refined the column's integrity by excluding 62 invalid samples.

Following this multi-stage cleaning protocol, we achieve a final dataset comprising 384,095 high-quality question-answer pairs, resulting in a substantial and reliable corpus tailored for developing language models in the Brazilian Portuguese medical domain.

3.3. Question-Type Annotation

To provide a fine-grained, semantic understanding of user intent within the corpus, we enriched the dataset with a synthetically generated feature column classifying the type of each user question. This structural annotation is critical for analyzing the dataset's composition (e.g., quantifying the types of user needs) and enables more targeted modeling applications. Following established practices in dataset construction (Alhuzali et al., 2024; Guo et al., 2018), we employed gpt-oss-120b (Agarwal et al., 2025) to annotate each user query. Through a structured prompt, available in Appendix A, the model assigned each question to one of the following seven predefined classes, which represent distinct facets of health-related inquiries:

- **Anatomy and Physiology:** Groups questions about the structure and function of body organs and systems.
- **Choosing Healthcare Professionals:** Captures doubts about which specialist to consult or how to access medical services.
- **Diagnosis:** Questions focused on the interpretation of symptoms, medical exams, or the identification of diseases.
- **Epidemiology:** Covers topics such as causes, risk factors, and disease complications.

- **Healthy Lifestyle:** Refers to questions about habits that influence health, such as diet, exercise, or sleep.
- **Treatment:** Encompasses questions about medications, therapies, side effects, and contraindications.
- **Other:** Includes questions that do not fit into the previous categories but are still related to health or the healthcare system.

4. Dataset Characterization

To characterize the final MedPT corpus and validate its suitability for training robust language models, we performed an extensive exploratory data analysis. This section details the dataset’s key statistical and structural properties, including an analysis of the distribution across question types and medical specialties, a summary of its linguistic features, and a presentation of representative samples. The analysis underscores the corpus’s large scale and multi-domain nature, confirming its value as a comprehensive resource for the Brazilian Portuguese medical NLP community.

4.1. Dataset Overview

The MedPT dataset is a large-scale corpus for medical question answering in Brazilian Portuguese. Each entry comprises a patient’s question paired with one or more verified answers from licensed healthcare professionals, ensuring both authenticity and clinical reliability. The dataset is structured with rich metadata to support a diverse range of downstream NLP tasks. The fields are detailed as follows:

- **Question:** The patient’s original inquiry in Brazilian Portuguese, typically describing symptoms, test results, or seeking information about medications and treatments.
- **Answer:** A response provided by a licensed healthcare professional, containing clinically grounded explanations or recommendations.
- **Condition:** A finer-grained topic within the main field, often a specific disease or condition (e.g., *Diabetes, Acne, Hypertension*).
- **Medical Specialty:** The specialization of the professional who authored the answer (e.g., *General Practitioner, Psychologist, Nutritionist*).
- **Question Type:** The semantic intent of the question, classified into one of seven categories: *Diagnosis, Treatment, Anatomy and Physiology, Epidemiology, Healthy Lifestyle, Choosing Healthcare Professionals, or Other*.

Condition	Questions	Answers
Frequent - Questions		
Syphilis	2,336	2,603
Inguinal Hernia	1,914	2,826
Appendicitis	1,773	2,250
Frequent - Answers		
Depression	1,140	11,715
Anxiety	1,025	10,734
Bipolar Disorder	928	3,906
Frequent - Both		
HIV and AIDS	5,664	6,162
COVID-19	4,674	5,339

Table 1: Most frequent conditions in the dataset, comprising the union of the five most frequent by questions and answers.

Collectively, these fields form a richly structured dataset comprising over 384K question-answer pairs, derived from more than 180K unique questions. This composition enables a wide range of applications, from specialty prediction and question-type classification to the development and domain adaptation of medical language models.

4.2. Quantitative Analysis

The scale and structure of MedPT are defined by its core statistics. A key characteristic of the dataset is its one-to-many mapping, with each question receiving an average of 2.11 answers. While most questions receive between one and three responses, some attract a much larger number, with an observed maximum of 266 answers for the same question. This multiplicity is a central feature of the corpus, reflecting the platform’s collaborative nature where different healthcare professionals provide complementary clinical perspectives on a single patient inquiry.

To characterize the linguistic properties of the corpus, we performed a quantitative analysis of its primary textual fields. Together, the question and answer columns contain nearly 57 million tokens distributed in a natural asymmetry that reflects the platform’s core pragmatic context. Patient questions are typically concise while the corresponding professional answers are substantially more detailed and explanatory (averaging 51.79 and 96.49 tokens length, respectively). This divergence stems from the distinct communicative goals of each role: patients are focused on symptom description and problem specification, whereas doctors are tasked with providing comprehensive explanations, differential considerations, and essential risk-management disclaimers.

Furthermore, the high standard deviations for to-

Condition	Questions	Answers
Food desensitization	1	39
Nightmares	1	38
Changes After Death	2	75
Relationship conflicts	3	112
Mutism	1	35
Existential distress	2	66
Difficulty in decision-making	3	96
Deep and prolonged sadness	3	95
Shyness	10	311

Table 2: Conditions with few patient questions and a high number of answers.

ken counts in both fields (37.5 for questions and 82.7 for answers) highlight a strong variability in linguistic depth. This variability is a key feature of the resource, demonstrating that the dataset captures a wide spectrum of interactions, from simple and direct inquiries to complex narrative-driven case descriptions, making it a robust resource for training models that can both understand diverse patient communication styles and generate appropriately detailed expert responses.

4.3. Diversity and Coverage

Beyond its scale, MedPT exhibits substantial thematic diversity, covering over 3,200 unique medical conditions. An analysis of the condition distribution reveals wide coverage across medical domains and highlights a key characteristic of the dataset. As shown in Table 1, the most frequent conditions among patient questions include infectious and clinical diseases such as HIV/AIDS, COVID-19, and Syphilis, while the conditions with the highest volume of professional answers are concentrated in mental health fields. This distinction underscores the dataset’s unique value, capturing a broad spectrum of both urgent public health concerns from patients and areas of deep, sustained engagement from healthcare professionals.

Furthermore, Table 2 demonstrates that conditions with few questions often received a huge amount of answers, suggesting that even rare or highly specific issues tend to stimulate extensive professional interaction. This pattern indicates that, although these conditions represent a smaller portion of the dataset, they attract multiple perspectives and deeper discussions around the same underlying question.

To facilitate a finer-grained analysis of user intent, each question was enriched with a *Question Type* label. The distribution of these categories, shown in Figure 1, reveals that the dataset is dominated by action-oriented inquiries. Questions concerning *Treatment* and *Diagnosis* collectively account




Figure 1: Distribution of medical questions by type.

for nearly 70% of the data, indicating a primary focus on actionable medical guidance. Concurrently, the presence of informational categories such as *Epidemiology* and *Healthy Lifestyle* confirms the dataset’s thematic breadth. This diversity is also reflected in the specialties of the answer providers, including answers from professionals across 104 distinct specialties, which form over 900 unique specialization combinations due to multi-specialty declarations.

Furthermore, the data exhibits a long-tail distribution characteristic of real-world: high-frequency specialties such as *Psychology*, *Gynecology*, and *Orthopedics* are well-represented (as shown in Figure 2), while a wide array of highly specialized fields, including *Bariatric Surgery* and *Pediatric Nutritionology*, constitute the long tail with few number of examples (8 and 4 samples, respectively). This broad spectrum, spanning both common health concerns and niche medical domains, underscores the dataset’s ecological validity as a resource for modeling authentic healthcare interactions, offering a broad and realistic coverage.

5. Experiments

To test MedPT dataset under downstream applications, we formulate a multi-class classification task designed to predict the appropriate medical specialty from a patient’s question. This task simulates a critical real-world use case: the automated routing of patient inquiries to the correct clinical department. To evaluate model robustness against

varying levels of label complexity, we create three experimental tiers based on the Top 5, Top 10, and Top 20 most frequent subspecialties in the dataset.




Figure 2: Distribution of the top five most represented medical specialties by number of answers.

5.1. Experimental Setup

For all experiments, we use the Qwen3 1.7B (Yang et al., 2025) due to its balance between computational efficiency and performance, offering competitive results while maintaining a substantially lower training and inference cost compared to larger-scale models. We then benchmark its performance across three distinct methodological settings corresponding to the Top 5, Top 10, and Top 20 most frequent subspecialties. Each split is divided into an 80% training set and a 20% held-out test set. The evaluation settings are as follows:

1. **Zero-Shot:** The base model performs classification using only the task instructions and the patient’s question, without any examples. This tests the model’s intrinsic ability to generalize to the medical domain.
2. **Few-Shot (In-Context Learning):** The base model is prompted with instructions along with a small number of dynamically-sampled examples from the training set. We evaluate its in-context learning capability with 3 and 5 examples to measure performance with varying levels of contextual information.
3. **Fine-Tuning:** The model is fine-tuned on the 80% training split of each respective sub-dataset. The training is conducted for 5 epochs

using the AdamW optimizer with a learning rate of $2e-5$ and a batch size of 32. This setting evaluates the dataset’s effectiveness for specializing a model to the task.

Performance across all settings is measured using standard classification metrics: precision, recall, and F1-score, calculated with macro averages to account for potential class imbalance.

	Strategy	Precision	Recall	F1-Score
Top-5	Zero-Shot	0.81	0.60	0.60
	3-Shot	0.88	0.86	0.86
	5-Shot	0.88	0.87	0.87
	Fine-Tuned	0.98	0.98	0.98
Top-10	Zero-Shot	0.79	0.55	0.54
	3-Shot	0.84	0.76	0.78
	5-Shot	0.85	0.77	0.79
	Fine-Tuned	0.96	0.96	0.96
Top-20	Zero-Shot	0.75	0.58	0.60
	3-Shot	0.80	0.71	0.73
	5-Shot	0.80	0.71	0.73
	Fine-Tuned	0.94	0.94	0.94

Table 3: Results for classification.

5.2. Results

The experimental results, summarized in Table 3, reveal a clear and consistent performance hierarchy across all tested configurations. Fine-tuning the Qwen3-1.7B model on the MedPT dataset yields the highest performance by a significant margin, establishing a strong benchmark and demonstrating the dataset’s value for model specialization. As expected, performance for all methods degrades systematically as the task difficulty increases from 5 to 20 classes. Nonetheless, the fine-tuned model maintains an exceptional 0.94 F1-score on the most challenging 20-class task, underscoring the quality of the learned representations.

The efficacy of in-context learning (ICL) is also clearly demonstrated. In the 20-class setup, for example, the zero-shot baseline achieves a 0.60 F1-score, indicating the task’s inherent difficulty. However, by providing just three examples (3-shot), the score jumps to 0.73. This substantial improvement highlights the model’s ability to quickly leverage the contextual patterns within MedPT to enhance its predictive accuracy without full fine-tuning.

A key finding from the ICL experiments is a clear point of diminishing returns. The performance gain from zero-shot to 3-shot is substantial across all settings. However, increasing the examples from 3-shot to 5-shot provides only a marginal benefit (e.g., 0.73 F1-score for both in the 20-class setup), suggesting that a 3-shot prompt may represent the optimal balance of performance and efficiency for this task.

6. Discussion

The experimental results confirm the high utility of MedPT for developing specialized medical language models. The exceptional performance of the fine-tuned model, achieving an F1-score of 0.94 on the medical specialty classification task involving 20 specialties, directly validates the dataset's quality. This demonstrates that the linguistic patterns and class labels within MedPT are consistent and highly learnable, providing a robust foundation for domain-specific classifiers. Furthermore, our findings offer practical implications for deployment: the substantial performance gains from in-context learning confirm MedPT's value for resource-light applications, while the observed diminishing returns between 3-shot and 5-shot prompts suggest an optimal efficiency-performance trade-off.

A deeper qualitative analysis of the fine-tuned model's errors revealed that misclassifications were not random but systematic, occurring between semantically and clinically related subspecialties. We observed notable confusion between conditions with high comorbidity (e.g., *Anxiety* and *Depression*), anatomical proximity (e.g., *Umbilical Hernia* and *Inguinal Hernia*), or overlapping symptoms (e.g., *Gastritis* and *Gastroesophageal Reflux*). This pattern indicates that the model is capturing a meaningful semantic space of medical concepts, where its errors reflect genuine clinical ambiguity, underscoring MedPT's potential to support future research into more complex tasks, such as fine-grained medical entity disambiguation.

7. Conclusion

In this work, we introduced MedPT, the first large-scale, real-world dataset of patient-doctor interactions in Brazilian Portuguese, containing over 384,000 curated question-answer pairs. This resource was developed to address the critical gap of culturally and clinically relevant data for a major but under-resourced language, capturing authentic user inquiries that include endemic diseases and local health concerns often absent in translated datasets. Our contribution extends beyond collection to the rigorous, multi-stage curation protocol itself. We applied a hybrid quantitative-qualitative analysis to every feature, using statistical distributions to guide deep manual inspection. This meticulous process, coupled with an LLM-driven annotation of user intent, ensured the dataset's high signal-to-noise ratio and structural richness.

This high utility was then confirmed through multifaceted validation. In downstream classification experiments, a 1.7B-parameter model fine-tuned on MedPT achieved an F1-score of 0.94 on the medical specialty classification task across 20 spe-

cialties, compared to 0.60 for the base model. This, combined with strong few-shot ICL results, proves the dataset's practical value. By publicly releasing MedPT, we provide a foundational resource to the research community. We aim to support the development of more equitable, accurate, and safe medical NLP technologies, enabling a new generation of models that can understand the specific clinical and linguistic nuances of the Portuguese-speaking world.

8. Limitations

We acknowledge several limitations in this work that also present clear avenues for future research. First, while our experimental validation focused exclusively on a multi-class classification task demonstrating the dataset's utility for a crucial application like inquiry routing, its performance on generative tasks, such as medical question-answering and patient-doctor dialogue, remains to be explored. Second, our deduplication process was based on removing exact-match samples instead of more sophisticated near-duplication detection techniques. Finally, regarding data quality, although all answers originate from verified physicians, a secondary, independent clinical review of the responses could further reinforce the dataset's reliability and ensure the clinical accuracy of the provided information.

9. Ethical Considerations

The MedPT dataset and any models developed using it are intended to function as assistive tools to support, not replace, the judgment of qualified healthcare professionals. We emphasize that AI systems trained on this data should not be used as a substitute for professional medical consultation, diagnosis, or treatment. Relying solely on LLM-generated content for medical decisions carries significant risks, as models may produce outputs that are plausible but clinically inaccurate. The primary ethical application of this work is to augment clinical workflows, for tasks such as inquiry routing, information retrieval, preliminary symptom summarization, while ensuring that the final diagnostic and therapeutic decisions remain the exclusive responsibility of a human medical expert.

10. Acknowledgments

This work has been fully/partially funded by the project Research and Development of Algorithms for Construction of Digital Human Technological Components supported by Advanced Knowledge Center in Immersive Technologies (AKCIT), with financial resources from the PPI IoT of the MCTI grant number 057/2023, signed with EMBRAPPII.

11. Bibliographical References

- Mohammed Abdelhay, Ammar Mohammed, and Hesham A Hefny. 2023. Deep learning for arabic healthcare: Medicalbot. *Social Network Analysis and Mining*, 13(1):71.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Al-tenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Hassan Alhuzali, Ashwag Alasmari, and Hamad Al-saleh. 2024. Mentalqa: An annotated arabic corpus for questions and answers of mental health-care. *IEEE Access*.
- Balsam Alkouz, Zaher Al Aghbari, and Jemal Hussien Abawajy. 2019. [Tweeluenza: Predicting flu trends from twitter data](#). *Big Data Mining and Analytics*, 2(4):273–287.
- Marco Cascella, Federico Semeraro, Jonathan Montomoli, Valentina Bellini, Ornella Piazza, and Elena Bignami. 2024. The breakthrough of large language models release for medical applications: 1-year timeline and perspectives. *Journal of Medical Systems*, 48(1):22.
- Bernardo Scapini Consoli, Henrique Dias P dos Santos, Ana Helena DPS Ulbrich, Renata Vieira, and Rafael Heitor Bordini. 2022. Brateca (brazilian tertiary care dataset): a clinical information dataset for the portuguese language. In *Proceedings of the 13th Conference on Language Resources and Evaluation (LREC 2022), Marseille, 20-25 June, 2022, França*.
- Ricardo C Cury, Istvan Megyeri, Tony Lindsey, Robson Macedo, Juan Battle, Shwan Kim, Brian Baker, Robert Harris, and Reese H Clark. 2021. Natural language processing and machine learning for detection of respiratory illness by chest ct imaging and tracking of covid-19 pandemic in the united states. *Radiology: Cardiothoracic Imaging*, 3(1):e200596.
- Fábio André da Costa Lopes. 2019. Contributions to clinical information extraction in portuguese: Corpora, named entity recognition, word embeddings. Master's thesis, Universidade de Coimbra (Portugal).
- Nastaran Enshaei and Farnoosh Naderkhani. 2024. [The role of data quality for reliable ai performance in medical applications](#). *IEEE Reliability Magazine*, 1(3):24–28.
- Fernanda B Färber, Julia S Dollis, Pedro SFB Ribeiro, Iago A Brito, Rafael T Sousa, and Arlindo R Galvão Filho. 2025. Iaramed: A women's healthcare chatbot for portuguese speakers. In *Simpósio Brasileiro de Computação Aplicada à Saúde (SBCAS)*, pages 931–942. SBC.
- Mehrdad Ghassabi, Pedram Rostami, Hamidreza Baradaran Kashani, Amirhossein Poursina, Zahra Kazemi, and Milad Tavakoli. 2025. Leveraging online data to enhance medical knowledge in a small persian language model. *arXiv preprint arXiv:2505.16000*.
- Haihong Guo, Xu Na, and Jiao Li. 2018. Qcorp: an annotated classification corpus of chinese health questions. *BMC medical informatics and decision making*, 18(Suppl 1):16.
- Junqing He, Mingming Fu, and Manshu Tu. 2019. Applying deep matching networks to chinese medical question answering: a study and a dataset. *BMC medical informatics and decision making*, 19(Suppl 2):52.
- Xuehai He, Shu Chen, Zeqian Ju, Xiangyu Dong, Hongchao Fang, Sicheng Wang, Yue Yang, Jiaqi Zeng, Ruisi Zhang, Ruoyu Zhang, et al. 2020. Meddialog: Two large-scale medical dialogue datasets. *arXiv preprint arXiv:2004.03329*.
- Bin Ji, Rui Liu, Shasha Li, Jie Yu, Qingbo Wu, Yuesong Tan, and Jiaju Wu. 2019. A hybrid approach for named entity recognition in chinese electronic medical record. *BMC medical informatics and decision making*, 19(Suppl 2):64.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*.
- Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, 10(1):1.
- Samir Khan, Seiji Tsutsumi, Takehisa Yairi, and Shinichi Nakasuka. 2021. [Robustness of ai-based prognostic and systems health management](#). *Annual Reviews in Control*, 51:130–152.
- Qiming Liu, Ruirong Yang, Qin Gao, Tengxiao Liang, Xiuyuan Wang, Shiju Li, Bingyin Lei, and Kaiye Gao. 2025. [A review of applying large](#)

- language models in healthcare. *IEEE Access*, 13:6878–6892.
- Ling Luo, Po-Ting Lai, Chih-Hsuan Wei, Cecilia N Arighi, and Zhiyong Lu. 2022. [BioRED: a rich biomedical relation extraction dataset](#). *Briefings in Bioinformatics*, 23(5):bbac282.
- Andrés Alejandro Ramos Magna, Héctor Allende-Cid, Carla Taramasco, Carlos Becerra, and Rosa L Figueroa. 2020. Application of machine learning and word embeddings in the classification of cancer diagnosis using patient anamnesis. *Ieee Access*, 8:106198–106213.
- Andriy Mulyar, Ozlem Uzuner, and Bridget McInnes. 2021. [Mt-clinical bert: scaling clinical information extraction with multitask learning](#). *Journal of the American Medical Informatics Association*, 28(10):2108–2115.
- Palanichamy Naveen and Pavel Trojovský. 2024. Overview and challenges of machine translation for contextually appropriate translations.” *iscience* 27 (10): 110878.
- Z Al Nazi and W Peng. 2024. Large language models in healthcare and medical domain: A review. *informatics*, 11, 57.
- Nhung Thi-Hong Nguyen, Phuong Phan-Dieu Ha, Luan Thanh Nguyen, Kiet Van Nguyen, and Ngan Luu-Thuy Nguyen. 2022. Spbertqa: A two-stage question answering system based on sentence transformers for medical texts. In *International Conference on Knowledge Science, Engineering and Management*, pages 371–382. Springer.
- Vincent Nguyen, Sarvnaz Karimi, Maciej Rybinski, and Zhenchang Xing. 2023. Medredqa for medical consumer question answering: Dataset, tasks, and neural baselines. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 629–648.
- Lucas Emanuel Silva e Oliveira, Ana Carolina Peters, Adalniza Moura Pucca Da Silva, Caroline Pilatti Gebelua, Yohan Bonescki Gumiel, Lilian Mie Mukai Cintho, Deborah Ribeiro Carvalho, Sadid Al Hasan, and Claudia Maria Cabral Moro. 2022. Semclinbr-a multi-institutional and multi-specialty semantically annotated corpus for portuguese clinical nlp tasks. *Journal of Biomedical Semantics*, 13(1):13.
- Pedro Henrique Paiola, Gabriel Lino Garcia, Joao Renato Ribeiro Manesco, Mateus Roder, Douglas Rodrigues, and João Paulo Papa. 2024. Adapting llms for the medical domain in portuguese: A study on fine-tuning and model evaluation. *arXiv preprint arXiv:2410.00163*.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.
- Sam Preston, Mu Wei, Rajesh Rao, Robert Tinn, Naoto Usuyama, Michael Lucas, Yu Gu, Roshanthi Weerasinghe, Soohye Lee, Brian Piening, et al. 2023. Toward structuring real-world data: Deep learning for extracting oncology information from clinical text with patient-level supervision. *Patterns*, 4(4).
- Philippa Ryan Conmy, Berk Ozturk, Tom Lawton, and Ibrahim Habli. 2023. The impact of training data shortfalls on safety of ai-based clinical decision support systems. In *Computer Safety, Reliability, and Security*, pages 213–226, Cham. Springer Nature Switzerland.
- Daniel Schwabe, Katinka Becker, Martin Seyferth, Andreas Klaß, and Tobias Schaeffter. 2024. The metric-framework for assessing data quality for trustworthy ai in medicine: a systematic review. *NPJ digital medicine*, 7(1):203.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2022. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R Pfohl, Heather Cole-Lewis, et al. 2025. Toward expert-level medical question answering with large language models. *Nature Medicine*, 31(3):943–950.
- Jing Wang, Huan Deng, Bangtao Liu, Anbin Hu, Jun Liang, Lingye Fan, Xu Zheng, Tong Wang, and Jianbo Lei. 2020. Systematic evaluation of research progress on natural language processing in medicine over the past 20 years: bibliometric study on pubmed. *Journal of medical Internet research*, 22(1):e16816.
- Stephen Wu, Kirk Roberts, Surabhi Datta, Jingcheng Du, Zongcheng Ji, Yuqi Si, Sarvesh Soni, Qiong Wang, Qiang Wei, Yang Xiang, Bo Zhao, and Hua Xu. 2019. [Deep learning in clinical natural language processing: a methodical review](#). *Journal of the American Medical Informatics Association*, 27(3):457–470.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Sheng Zhang, Xin Zhang, Hui Wang, Lixiang Guo, and Shanshan Liu. 2018. [Multi-scale attentive interaction networks for chinese medical question answer selection](#). *IEEE Access*, 6:74061–74071.

Enwei Zhu, Qilin Sheng, Huanwan Yang, Yiyang Liu, Ting Cai, and Jinpeng Li. 2023. [A unified framework of medical information annotation and extraction for chinese clinical text](#). *Artificial Intelligence in Medicine*, 142:102573.

Mert İncidelen and Murat Aydoğın. 2024. [Developing question-answering models in low-resource languages: A case study on turkish medical texts using transformer-based approaches](#). In *2024 8th International Artificial Intelligence and Data Processing Symposium (IDAP)*, pages 1–4.

A. Question-Type Annotation Prompt

The following prompt (originally written in Brazilian Portuguese) was used to instruct `gpt-oss-120b` to classify each patient question into one of seven predefined semantic categories, as described in Section 3.3.

Prompt

Act as a classifier specialized in medicine, with deep knowledge of all types of medical questions. Your task is to classify medical questions based solely on the content of the question.

Instructions:

- (1) Classify the question into only one of the types listed below.
- (2) Never classify into more than one type.
- (3) If it is not possible to classify with confidence, classify as “Others”.
- (4) Your answer must contain only the TYPE of the question. Nothing else.
- (5) Do not invent types that are not listed.

Available types:

Diagnosis, Treatment, Anatomy and Physiology, Epidemiology, Healthy Lifestyle, Choosing Healthcare Professionals, Others.

Data to be classified:

```
-- Start of data --  
Question: {question}  
-- End of data --
```