

Enhanced medical data extraction: leveraging LLMs for accurate retrieval of patient information from medical reports

Angel Manuel Garcia-Carmona, Maria-Lorena Prieto, Enrique Puertas, Juan-Jose Beunza

Submitted to: JMIR AI
on: November 14, 2024

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript.....	5
---------------------------------	----------

Preprint
JMIR Publications

Enhanced medical data extraction: leveraging LLMs for accurate retrieval of patient information from medical reports

Angel Manuel Garcia-Carmona¹; Maria-Lorena Prieto¹; Enrique Puertas^{1, 2, 3} PhD; Juan-Jose Beunza^{1, 4, 5, 2} MD, PhD

¹ Research and Doctorate School Universidad Europea de Madrid Madrid ES

² IASalud Universidad Europea de Madrid Madrid ES

³ Department of Computing and Technology Univesidad Europea de Madrid Madrid ES

⁴ Hospital La Paz Institute for Health Research – IdiPAZ (Universidad Europea de Madrid) Madrid ES

⁵ Department of Medicine Universidad Europea de Madrid Madrid ES

Corresponding Author:

Juan-Jose Beunza MD, PhD

Research and Doctorate School
Universidad Europea de Madrid
Calle Tajo s/n
Madrid
ES

Abstract

Background: The digital transformation of healthcare has introduced both opportunities and challenges, particularly in managing and analyzing the vast amounts of unstructured medical data generated daily. There is a need to explore the feasibility of generative solutions in extracting data from medical reports, categorized by specific criteria.

Objective: This study investigates the application of Large Language Models (LLMs) for the automated extraction of structured information from unstructured medical reports, employing the LangChain framework in Python.

Methods: Through a systematic evaluation of leading LLMs—GPT-4o, LLaMA 3, LLaMA 3.1, Gemma 2, Qwen 2, and Qwen 2.5—using zero-shot prompting techniques and embedding results into a vector database, the research assesses their performance in extracting patient demographics, diagnostic details, and pharmacological data.

Results: Evaluation metrics, including accuracy, precision, recall, and F1 scores, revealed high efficacy across most categories, with GPT-4o achieving the highest overall performance (91.4% accuracy).

Conclusions: The findings highlight notable differences in precision and recall between models, particularly in extracting names and age-related information. Challenges in processing unstructured medical text, including variability in model performance across data types, are discussed. The study demonstrates the feasibility of integrating LLMs into healthcare workflows, offering significant improvements in data accessibility and supporting clinical decision-making processes. Additionally, it explores the role of retrieval-augmented generation (RAG) techniques in enhancing information retrieval accuracy, addressing issues such as hallucinations and outdated data in LLM outputs. Future work emphasizes the need for optimization through larger and more diverse training datasets, advanced prompting strategies, and the integration of domain-specific knowledge to improve model generalizability and precision.

(JMIR Preprints 14/11/2024:68776)

DOI: <https://doi.org/10.2196/preprints.68776>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

✓ **Please make my preprint PDF available to anyone at any time (recommended).**

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to the public.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <https://www.jmir.org/2018/01/e10000/>

No. Please do not make my accepted manuscript PDF available to anyone. I understand that if I later pay to participate in <https://www.jmir.org/2018/01/e10000/>



Original Manuscript

1 **Article**2 **Enhanced medical data extraction: leveraging**
3 **LLMs for accurate retrieval of patient**
4 **information from medical reports**5 **Ángel M García-Carmona¹, Maria-Lorena Prieto¹, Enrique Puertas ^{1,2,3}, Juan-Jose**
6 **Beunza^{1,3,4,5*}**

Citation: To be added by
editorial staff during production.

Academic Editor: Firstname
Lastname

Received: date

Revised: date

Accepted: date

Published: date



Copyright: © 2024 by the
authors. Submitted for possible
open access publication under
the terms and conditions of the
Creative Commons Attribution
(CC BY) license
(<https://creativecommons.org/licenses/by/4.0/>).

* Correspondence: juanjose.beunza@universidadeuropea.es

¹ Research and Doctorate School, Universidad Europea de Madrid, Madrid, Spain.

² Department of Computing and Technology, Universidad Europea de Madrid, Madrid, Spain.

³ IASalud, Universidad Europea de Madrid, Madrid, Spain.

⁴ Hospital La Paz Institute for Health Research – IdiPAZ (Universidad Europea de Madrid), Madrid, Spain.

⁵ Department of Medicine, Universidad Europea de Madrid, Madrid, Spain.

Angel M. García-Carmona: 21951211@live.uem.es

Maria-Lorena Prieto: 20811530@live.uem.es

Enrique Puertas: enrique.puertas@universidadeuropea.es

Juan-Jose Beunza: juanjose.beunza@universidadeuropea.es

Background: The digital transformation of healthcare has introduced both opportunities and challenges, particularly in managing and analyzing the vast amounts of unstructured medical data generated daily. There is a need to explore the feasibility of generative solutions in extracting data from medical reports, categorized by specific criteria.

Objective: This study investigates the application of Large Language Models (LLMs) for the automated extraction of structured information from unstructured medical reports, employing the LangChain framework in Python.

Methods: Through a systematic evaluation of leading LLMs—GPT-4o, LLaMA 3, LLaMA 3.1, Gemma 2, Qwen 2, and Qwen 2.5—using zero-shot prompting techniques and embedding results into a vector database, the research

assesses their performance in extracting patient demographics, diagnostic details, and pharmacological data.

Results: Evaluation metrics, including accuracy, precision, recall, and F1 scores, revealed high efficacy across most categories, with GPT-4o achieving the highest overall performance (91.4% accuracy).

Conclusions: The findings highlight notable differences in precision and recall between models, particularly in extracting names and age-related information. Challenges in processing unstructured medical text, including variability in model performance across data types, are discussed. The study demonstrates the feasibility of integrating LLMs into healthcare workflows, offering significant improvements in data accessibility and supporting clinical decision-making processes. Additionally, it explores the role of retrieval-augmented generation (RAG) techniques in enhancing information retrieval accuracy, addressing issues such as hallucinations and outdated data in LLM outputs. Future work emphasizes the need for optimization through larger and more diverse training datasets, advanced prompting strategies, and the integration of domain-specific knowledge to improve model generalizability and precision.

Keywords: Large Language Models, Langchain framework, Electronic Health Records, Data mining, Model Evaluation, Healthcare, Digitalization

Introduction

In recent years, the healthcare sector has witnessed a significant shift towards digital systems for managing medical information, including electronic health records, diagnostic imaging tests, and bureaucratic records. This transition has been further accelerated by the COVID-19 pandemic, which popularized telemedicine as a means to reduce contagion risks, minimize travel, and improve access to healthcare in remote areas [1-3]. However, the increasing reliance on digital systems has led to the generation of vast amounts of unstructured medical data, posing challenges for efficient information extraction and utilization.

The complexity of managing unstructured medical data necessitates innovative approaches to support clinical and sociological studies, optimize research, and enhance diagnostic precision. In this context, the potential of generative Artificial Intelligence (AI) solutions, particularly Large Language Models (LLMs), has emerged as a promising avenue for automating the extraction of structured clinical information from unstructured medical reports.

This study investigates the feasibility of leveraging LLMs, specifically through the LangChain framework, to address key challenges in healthcare data digitalization, such as accuracy, scalability, and integration into existing workflows. It evaluates the performance of leading LLMs in extracting critical data categories, including patient demographics, diagnostic details, and pharmacological information. By exploring the capabilities of generative AI in this domain, the research aims to enhance clinical decision-making, optimize resource allocation, and improve overall efficiency in healthcare systems.

Background

Large Language Models (LLMs) are advanced artificial intelligence systems grounded in deep learning architectures, predominantly utilizing transformer networks, that are trained on extensive textual corpora. These models are designed to capture complex linguistic patterns and semantic relationships, enabling them to process, generate, and predict human language with a high degree of accuracy. In healthcare, LLMs have the potential to contribute to transformative changes in healthcare by improving diagnostic precision, assisting in clinical decision-making processes, and facilitating communication between patients and healthcare providers [4,5] LLMs are capable of delivering foundational knowledge, contextual analysis, and accessible information, making them valuable tools for patient education and clinical consultations [6]. They can also be integrated into medical practice responsibly and effectively, providing tools that address the needs of various medical disciplines and diverse patient populations [7].

LLMs can be pre-trained models, meaning they possess the capacity to comprehend and generate text without the need for extensive additional training. This capability introduces significant challenges in managing the vast quantities of unstructured medical data generated, as extracting relevant information from these sources is inherently complex. LLMs, employing a transformer architecture, excel in a multitude of domains, demonstrating remarkable capabilities in natural language processing tasks and text comprehension. The essence of pre-training lies in enabling these models to predict the next word in a given text, a computational process that underpins their performance across various tasks, demonstrating their advanced design [8]. Transformers are based on multi-layer neural networks which are trained with large datasets.

Traditionally, text processing has been conducted using recurrent neural networks (RNNs), a type of neural network architecture specifically designed to handle sequential data, such as text. RNNs eliminate the need for explicit word history modeling by naturally incorporating temporal dimensions, allowing the network to retain relevant information from previous time steps. RNNs operate by encoding feature vectors for each word, constructing input vectors from word embeddings, and incorporating outputs from prior hidden states, either through copying or time-step delays. Typically, the softmax function serves as the activation function. An important aspect of RNNs is backpropagation through time, which adjusts weights based on the sequence's context. Ultimately, the output layer produces a probability distribution for each word based on prior words and contextual features.

To enhance word prediction accuracy, the utilization of socio-linguistic features, such as sequences of discourse-related tags that provide syntactic information, was explored to enhance word prediction accuracy. Additionally, clustering techniques were used to delineate conversation topics, acknowledging that linguistic choices are influenced by the thematic context, while incorporating log-scaled frequency considerations. Furthermore, we factored in the socio-situational context, which encompasses variables such as the conversational context (e.g., interview, spontaneous discussion, phone call, or academic seminar), the relationships between participants, and their quantity. These considerations collectively contribute to a more precise word prediction model [9].

To equip Large Language Models (LLMs) for tackling intricate challenges and transcending the constraints of generalized composition inherent in thought chain prompts, which are often based

on limited examples, a novel "from more to less" prompting approach has been introduced. This innovative methodology aims to combine structured natural language processing techniques with self-consistent decoding mechanisms. The proposed approach unfolds in sequential phases, commencing with the decomposition and resolution of subproblems. This involves furnishing consistent examples showcasing the resolution of subproblems and compiling lists of previously answered subquestions along with their solutions. It is noteworthy to emphasize that consistent decoding, in this context, refers to the coherent and logical interpretation of information during the model's generation process. This 'from more to less' approach lays the foundation for leveraging bidirectional interactions, thereby enhancing the performance of LLMs in complex tasks [10].

Given the distinct characteristics of LLMs and specific operational considerations, our primary focus lies in addressing the challenges associated with healthcare digitalization. Our research places a significant emphasis on information extraction, with a notable shift towards document analysis as opposed to the conventional Extract, Transform, Load (ETL) or Extract, Load, Transform (ELT) processes commonly applied to structured datasets. This approach, broadly categorized, aims to unveil structured information from unstructured or semi-structured texts, representing a more expressive method that enhances communication.

It is essential to note that ETL typically refers to the process of extracting, transforming, and loading data into a structured format, while ELT reverses the sequence by loading data first and then transforming it. Our work underscores the significance of document analysis as a specialized area within the broader field of empirical Natural Language Processing (NLP), involving the extraction and encoding of information in the context of healthcare digitalization [11].

To be more precise, in our experiment, we will collect various medical reports in PDF format. Using prompts, we will attempt to extract diverse clinical information, such as age, weight, family medical history, date of birth, or potential allergies. This information will be used to enhance our local data model, thus optimizing diagnostic monitoring by reducing the need for manual inquiries and the time spent on non-automated searches. The experiment will be conducted through the implementation of the Langchain framework in Python, with concurrent utilization of models from OpenAI (GPT 4o), Meta (LLaMa 3 and 3.1), Google (Gemma and Gemma 2) and Alibaba (Qwen 2 and Qwen 2.5).

The advancements in LLMs have significantly expanded their applications across various domains, particularly in healthcare, where they have demonstrated substantial utility. As we explore the intricacies of LLMs, the transition from foundational understanding to practical implementation becomes evident. In the preceding chapter, an exploration of the fundamental architecture of LLMs underscores their training methods and transformative capacities across diverse disciplines.

Now, directing our attention to the practical implications of LLMs in healthcare digitalization, the integration of these models into medical practice is a critical area of focus. The ensuing discussion delves into the strategic application of LLMs to address intricate healthcare challenges, emphasizing their pivotal role in information extraction and meticulous document analysis. This discussion lays the groundwork for our empirical endeavors, where LLMs are employed to extract critical clinical information from medical reports, enabling the

optimization of diagnostic monitoring and reducing reliance on manual efforts.

Related Works

Multiple experiments have been conducted using LLMs to analyze documents, employing metrics that evaluate fluency (whether the generated text is coherent), correctness (if the prompt response is appropriate), and the quality of citations (if the cited passages are suitable). These experiments involve combining automated metrics with human evaluation, which uses qualitative metrics to assess aspects such as utility and the coherence of citations, assigning scores on a scale from 1 to 5. Evaluation metrics were adapted for each dataset, incorporating custom accuracy measures tailored to the specifics of each dataset [12].

Network syntactic analysis, a method utilized for modeling knowledge about document components by delineating their geometric properties, lexical entities, and relationships, has emerged as a prominent technique. An example of its application is seen in the utilization of the FRESCO semantic network language. In this experiment, FRESCO was employed to analyze business letters, facilitating the extraction of structural elements such as the sender, recipient, date, and main body. This approach facilitates a comprehensive specification of knowledge concerning these structural components, contributing to accuracy and completeness in the modeling process.

The accuracy of structural entity recognition is high when the visual organization of document elements (position, size, images, text formatting, etc.) can be used to identify the sender. However, this accuracy may decrease when the information is not concentrated in a specific location and is instead scattered across different sections of the letter. This situation can lead to document rejection, but the use of network analysis, combined with layer-specific knowledge, can optimize information extraction and automatic response generation [13].

Given that current transformer-based neural networks employ probabilistic techniques, it is interesting to note that decades ago, probabilistic experiments were conducted based on research into the use of logistic regression for obtaining ad hoc data, where a regression equation is fitted to learn data. The variables used in the equation are often statistical averages. Linear regression is used to identify simple yet effective probabilistic paths by combining search cues. The effectiveness of information retrieval has been enhanced through manual reformulations of topics [14].

This approach mirrors earlier probabilistic retrieval methods, such as staged logistic regression (SLR), which combined multiple retrieval clues to improve relevance estimates. These foundational techniques, though simpler, share conceptual similarities with modern transformer-based models, where embeddings and attention mechanisms probabilistically weigh token importance. The evolution from manual reformulations and regression-based methods to automated neural networks underscores the enduring role of probabilistic thinking in enhancing retrieval effectiveness. [28]

The integration of knowledge graph structures has emerged as a pivotal resource in the realm of text document analysis. Through the

application of advanced natural language processing techniques, this approach facilitates the extraction of critical entities, such as geographical locations, temporal references, and personal names, followed by the utilization of specialized tools to address ambiguities and spelling variations. This approach, known as 'occurrence data,' emphasizes the preservation of terms, phrases, and entities throughout the analytical process.

The integration of natural language processing (NLP) with knowledge graph structures enhances textual comprehension by focusing on contextual relationships, which facilitates precise information retrieval and analysis. The use of knowledge graph structures in text document analysis enables deeper insights and a refined understanding of data, overcoming the limitations of traditional keyword-based search approaches and expanding the scope of scientific exploration and data analysis.

Once the various entities have been extracted, the knowledge graph construction process commences. Each extracted entity represents a labeled node, and for each source of the various entities, a corresponding node is added. In the graph, a weight-1 edge is introduced between entities that co-occur within a document, signifying their simultaneous presence. However, when adding new nodes to the graph, care must be taken to ensure that no pre-existing node with the label of the entity already exists, as in such cases, the existing node is repurposed.

In order to account for the diverse nature of entities, each node is equipped with a set of nature properties, allowing us to record the type of entities (e.g., distinguishing between individuals and geolocations). If a vertex to be inserted already exists, as is the case with locations and dates, the vertex's weight is increased incrementally. The resulting graph is both weighted and undirected, offering a wide range of query capabilities that can be tailored as needed. The structure of the links between nodes also allows for flexibility in the types of data that can be retrieved [15].

Recent advancements in knowledge graph (KG) augmentation have demonstrated the benefits of integrating textual information to enhance entity representations. For instance, recent work [29] proposes a multi-task framework that leverages dense retrieval to select highly relevant text descriptions for KG entities, subsequently augmenting the KG embeddings with these descriptions. This approach addresses the limitations of using single text descriptions by introducing a retriever model that automatically identifies richer and more contextually relevant text sources. Building on these advancements, this study explores the application of Graph Neural Networks (GNNs) in Natural Language Processing (NLP), focusing on Knowledge Graph Rewiring and Document Classification.

Leveraging GNNs' capabilities, we advance text analysis by uncovering hidden semantic connections and improving recommendation systems. By employing GNN-driven techniques to analyze semantic graphs and detect complex patterns in text data, our comparative analysis of GNN models, applied to knowledge graphs derived from modern art biographies, demonstrates their potential to enhance classification accuracy, manage noise, and provide deeper insights into text construction. These methods, combined with transformer-based models like SBERT for encoding text descriptions, achieve significant performance gains, highlighting the importance of integrating multiple text descriptions to capture diverse contexts. This research paves the way for broader applications of GNNs and dense

retrieval techniques in fields requiring detailed text analysis and sophisticated knowledge graph interpretation [30].

In the medical field, a substantial portion of data remains unstructured today, encompassing concepts such as emails, data streams, voice and video recordings, as well as digital documents. Structured data's growth tends to be more gradual. Automated text mining includes a range of methods that facilitate access to relevant information. Recent attention has been focused on natural language processing, as techniques from other domains, such as information retrieval and extraction (automated extraction of structured data from unstructured sources), are adapted and integrated into this context [16].

Data extraction often leads to the discovery of tabular data, which is frequently embedded within text, particularly in medical diagnoses. Traditional machine learning models struggle to efficiently process information in this format, while large language models (LLMs) also face limitations in this regard. In response, methodologies like TEMED-LLM have been developed, which include three key components: reasoning-extraction, result validation and correction, and training (preferably, of an interpretable model based on the extracted tabular data) [17].

With the aforementioned goal in mind, efforts have been directed towards tasks such as SCHEMA-TO-JSON, a task focused on the extraction of structured records from tables and other semi-structured data sources, such as a web page. This task takes as input a table that can optionally be supplemented with context from the same document, along with an extraction schema that specifies the attributes to be extracted for different records that may contain varying numbers of attributes. As a result, it generates a sequence of JSON objects represented by an array of key-value pairs, each paired with a record type, condensing the information into a more accessible format.

An approach for table extraction called InstructTE is applied, which demonstrates competitive performance in both accuracy and precision, with an emphasis on balancing the two. It only requires a human-constructed extraction schema, incorporating an error recovery strategy. The schema approach helps the extraction process adhere to a predefined structure, improving the accuracy and consistency of the extracted information. Primarily, human-driven prompting is employed to direct language models during the extraction of data from complex tables [18].

Additionally, other data extraction experiments have been conducted, focusing on radiological results that may not necessarily be textual reports. In the case of textual data, a state-of-the-art question-answering system was employed, contrasting with radiologist annotations [19]. On the other hand, for non-textual data, a manual extraction of various tomographies was performed, where the reports were randomly partitioned into training and validation sets based on a natural language rule to extract report attributes (resulting in high precision in identifying occlusion, distal, or basilar, of several large blood vessels) [20].

Materials and Methods

Models

In this study, we carefully selected specific versions and configurations of Large Language Models (LLMs) to ensure clarity and replicability in our experimental setup. For GPT-4o, the model used corresponds to the GPT-4o-2024-08-06 version, released in May 2024. This version, also known as GPT-4 Omni, is optimized for high-complexity, multi-step tasks, with training data extending until December 2023. GPT-4o includes a context window of up to 128,000 tokens and shows superior performance compared to GPT-4 Turbo, achieving twice the processing speed while reducing computational costs by 50%.

The LLaMA 3 model, developed and publicly released by Meta in April 2024, was evaluated in its 8-billion parameter (8B) configuration. This version incorporates a tokenizer vocabulary of 128,000 tokens and employs grouped query attention mechanisms to enhance performance on complex text tasks. Pretraining for LLaMA 3 was conducted on a dataset comprising 15 trillion tokens, with approximately 5% of the dataset consisting of languages other than English. Post-training included strategies such as supervised fine-tuning, preference optimization, rejection sampling, and proximal policy optimization. In addition to LLaMA 3, the updated LLaMA 3.1 version, released in July 2024, was also included in our study. LLaMA 3.1 incorporates architectural refinements, including enhanced attention mechanisms, and supports up to 405 billion parameters in its largest configuration. For this study, we utilized the 8B version of both models for consistency. Quantization from 16-bit to 8-bit numerics was applied to optimize computational performance, and both versions support a context window of up to 128,000 tokens.

The Qwen 2 and Qwen 2.5 models, developed by Alibaba Cloud, were evaluated in their 7-billion parameter (7B) configurations. Qwen 2, updated 3 months prior to this study, incorporates specialized bias terms for queries, keys, and values, significantly improving its attention mechanisms. This model was trained on a multilingual dataset spanning 27 languages, making it particularly robust for cross-linguistic applications. Qwen 2.5, released 2 months before our experiments, includes additional advancements in reasoning capabilities, such as Chain-of-Thought (CoT) and Program-of-Thought (PoT) techniques, which improve performance on tasks requiring structured and logical reasoning. Qwen 2.5 was pretrained on an expanded dataset of 18 trillion tokens, further refining its multilingual and contextual generation capabilities.

Finally, the Gemma models, derived from Google's Gemini generative chatbot, were also evaluated. Gemma 1, featuring 7 billion parameters, employs a decoder-only architecture designed for sequential text generation. It was trained with a context length of 8,192 tokens, optimizing it for tasks requiring strict sequence fidelity. Gemma 2, with 9 billion parameters, incorporates advanced techniques such as grouped-query attention and Root Mean Square Normalization (RMSNorm) to enhance its multi-head attention efficiency and model stability. Gemma 2 is particularly effective at selectively processing broader contexts while maintaining focus on smaller windows of words.

These configurations reflect a balance between computational feasibility and robust benchmarking across diverse model architectures. All models were evaluated under identical experimental conditions to ensure consistency and comparability of results.

Data

The documents under examination consist of clinical histories from diverse origins, lacking a standardized format. Sourced from various hospitals and medical conventions with heterogeneous organizational structures, these documents pose a unique challenge due to their non-conformity to a single medical specialty (e.g., cardiology, gynecology, psychiatry, etc.). This diversity results in a broad clinical and pharmacological spectrum, encompassing a wide range of clinical conditions and medication types.

The dataset used in this study comprises 100 Spanish medical reports in PDF format, carefully selected to represent a broad spectrum of clinical scenarios. These documents are unstructured medical records, primarily consisting of free-text narratives without a standardized format. They include sections related to patient demographics (e.g., age), clinical diagnoses, prescribed medications, diagnostic tests, and reasons for consultation. The length of the documents varies, with some being concise summaries and others containing more detailed descriptions of patient histories and treatments.

The heterogeneity of these cases, spanning various medical specialties (e.g., cardiology, internal medicine, family medicine) and levels of complexity, reflects the real-world variability encountered in clinical practice. This diversity is intentional, as the study aims to evaluate how effectively healthcare professionals can retrieve critical information from medical histories in time-sensitive clinical settings.

The dataset is fully anonymized, with no personally identifiable information included. The anonymization process was conducted by the source institutions (the Spanish Society of Internal Medicine, the Asturian Society of Family and Community Medicine, the Spanish Society of Cardiology, and the Faculty of Medicine at Francisco Marroquín University) prior to their provision for this study. These institutions followed their internal guidelines and ethical standards to ensure that all personal identifiers, such as patient names, addresses, and contact information, were removed or replaced with generic placeholders (e.g., 'Patient X'). This pre-existing anonymization ensures that the dataset is ethically compliant and suitable for research purposes.

While the lack of a standardized format poses challenges for information extraction, it also provides a realistic representation of the variability found in real-world medical records. This makes the dataset particularly valuable for evaluating the adaptability and robustness of Large Language Models (LLMs) in processing unstructured clinical data.

Spanish was chosen as the language for this study because, despite being the second most spoken language in the world, there is a noticeable gap in the number of studies conducted in Spanish compared to those in English. Addressing this gap is crucial to ensure that advancements in medical data processing technologies are accessible and applicable to Spanish-speaking healthcare professionals and systems. This focus enhances the study's relevance to a global audience while supporting the development of tools tailored to underrepresented linguistic contexts.

To extract data from these documents, we employ a zero-shot prompting data extraction technique [21], designed to enhance performance in tasks involving reasoning with linguistically untrained or previously unexposed information within a specific task or domain,

utilizing Pydantic. Building on this approach, we create a predefined prompt based on a template querying specific categories: "nombre" (name and surname), "edad" (age), "diagnostico" (diagnosis), "medicamentos" (drugs), and "pruebas" (medical tests). The prompt is structured as follows:

```
prompt_template = """
En base al contexto dado, responde en el formato indicado

Contexto: {context}

Pregunta: {question}

Formato de las instrucciones: {format_instructions}

Respuesta útil: """

QA_CHAIN_PROMPT = PromptTemplate(
    template=prompt_template,
    input_variables=["question", "context"],
    partial_variables={"format_instructions":
parser.get_format_instructions()},
)

rag_chain = (
    {"context": retriever, "question": RunnablePassthrough()}
    | QA_CHAIN_PROMPT
    | llm
    | parser
)

resultado = rag_chain.invoke("Obtener información sobre el
paciente").json(ensure_ascii=False)
```

This code defines the prompt template and the RAG (Retrieval-Augmented Generation) chain, which guides the model in extracting structured information from unstructured PDFs. The prompt is designed to ensure consistency and alignment with the predefined JSON schema. Using the information extracted from the documents, the code formats and prepares the corresponding prompt, proposing a structure that chains the categories according to predefined fields. This process ensures that the extracted data is organized and ready for further analysis or integration into downstream applications. It is important to note that our dataset comprises approximately 100 PDF documents written in the Spanish language.

Computational Resources and Implementation Details

This study was conducted using a personal computer equipped with an Intel Core i7 processor, an NVIDIA GeForce RTX graphics card, and 16 GB of RAM. This setup provided sufficient computational power to process the dataset and run the models efficiently within a local environment. While not utilizing extensive GPU clusters, this configuration demonstrates the feasibility of applying these methods using accessible hardware.

Preprocessing steps included extracting text from PDF files and segmenting the content into manageable chunks using a semantic chunker. The semantic chunker was specifically employed to ensure that the chunks maintained semantic coherence, a critical requirement to minimize hallucinations during information extraction, particularly in the sensitive context of healthcare. This approach allowed the

model to process contextually relevant pieces of information, thereby improving the reliability of the results. The processed chunks were stored in a FAISS vector database for retrieval purposes, although this specific choice of database did not influence model performance directly and was used primarily for organizational convenience.

The JSON schema was defined using Python's Pydantic library to ensure consistency in the extracted information. Prompt templates were carefully designed to query specific attributes, including name, age, diagnosis, medications, and tests, enabling structured data extraction.

Although the dataset itself cannot be shared due to confidentiality constraints, future work will explore the creation of synthetic datasets that mimic the structure and complexity of the original data to facilitate reproducibility. The implementation scripts used for processing, running models, and generating results are available upon reasonable request to the corresponding author. Detailed configurations, including prompt templates and hyperparameter settings, can also be shared to support replication efforts.

Retrieval Augmented Generation (RAG)

Retrieval Augmented Generation (RAG) is an approach that enhances language models by integrating information retrieval during the generation process, aiming to address issues like factual inaccuracies and hallucinations observed in the output of large language models [22]. The use of a semantic chunker ensured that only meaningful and contextually relevant information was fed into the retrieval and generation process, directly impacting the accuracy and reliability of the outputs. This methodological choice reflects the critical need for precision in healthcare applications, where even minor inaccuracies could lead to significant risks.

Traditional models, such as Naive RAG, follow a conventional methodology involving indexing, retrieval, and generation. In this paradigm, original data undergoes cleansing, conversion, and segmentation into manageable chunks represented as vectors through an embedding model. While Naive RAG provides a structured approach, it often faces challenges in retrieval precision, recall, and handling outdated information, which can affect the quality of generation. In this study, semantic chunking was utilized to address these challenges by ensuring that retrieved information retained contextual relevance, thus improving the reliability of the generation process in a critical domain like healthcare [23].

The Advanced RAG paradigm introduces optimization strategies in the pre-retrieval process, focusing on enhancing data indexing, fine-tuning embedding models, and post-retrieval processes such as re-ranking and prompt compression. Furthermore, the Modular RAG paradigm provides versatility and flexibility by integrating various methods to enhance functional modules, making it increasingly prevalent in the domain. Advanced RAG is considered a specialized form of Modular RAG, showcasing a relationship of inheritance and development among the three paradigms [24]. A sort of schematic graphic abstraction is shown in Fig. 1.




Figure 1: General schematic representation of Retrieval Augmented Generation for data extraction from documents.

While our exploration focused on the basic yet impactful facets of RAG, we specifically utilized a zero-shot prompting strategy combined with semantic chunking. The semantic chunker, implemented using LangChain, divides the extracted text into semantically coherent segments by analyzing the differences in embeddings between sentences. In our implementation, we used the SemanticChunker class with only the embeddings parameter configured, leveraging OpenAI embeddings to generate vector representations of the text. This approach ensures that the text is split into meaningful and contextually relevant chunks, which are then used in the retrieval and generation process.

The semantic chunker works by determining when to "break" apart sentences based on differences in their embeddings. When the difference between two sentences exceeds a predefined threshold (automatically calculated by the chunker), they are split into separate chunks. By relying solely on the embeddings parameter, we allowed the chunker to use its default settings for threshold calculation and chunk size, ensuring a balance between semantic coherence and practical usability. This simplicity in configuration was chosen to maintain efficiency while still achieving high-quality chunking results.

The retrieved chunks are integrated into the prompt as contextual information for the language model during response generation. To ensure the model prioritizes the retrieved information over its pre-trained knowledge, we structured the prompt to explicitly instruct the model to base its responses on the provided context. The prompt template included the context (the top 3 most relevant chunks, selected based on their embedding similarity to the query using FAISS), the user query or task to be performed, and format instructions to ensure the output adhered to the required structure.

The integration was implemented using LangChain, where the retriever (FAISS) fetched the most relevant chunks, and the prompt template combined these chunks with the query and format instructions. This approach ensured that the model's outputs were grounded in the retrieved information, aligning with the structured schema of the task. For a detailed implementation of the prompt structure and retrieval process, please refer to the Data chapter.

The semantic chunker, implemented using LangChain and the corresponding LLM embeddings, ensured that the retrieved chunks were semantically coherent and contextually relevant. By relying on embedding similarity (cosine similarity) and selecting the top 3 chunks, we minimized the risk of irrelevant or fragmented information being included in the prompt. The retrieval process was implemented using FAISS (Facebook AI Similarity Search), a highly efficient library for similarity search in high-dimensional spaces. FAISS indexed the document chunks as vector embeddings, enabling fast and accurate retrieval of the most relevant chunks based on their semantic similarity to the query.

When a query is received, its embedding is generated using the same embedding model employed for the document chunks. This ensures that both the query and the chunks are represented in the same vector space, allowing for a direct comparison of their semantic similarity. The similarity between the query embedding and each chunk embedding is calculated using cosine similarity, a metric that measures the cosine of the angle between two vectors in the embedding space. The top 3 chunks with the highest cosine similarity scores were selected for inclusion in the prompt.

This method inherently validated the relevance of the chunks, as only those with the highest similarity to the query were used. The decision to employ semantic chunking was informed by prior studies, which highlight its ability to enhance precision in information extraction and semantic analysis, particularly in complex domains like healthcare [31, 32]. The combination of FAISS for efficient retrieval and cosine similarity for semantic comparison ensured that the retrieved information was both accurate and contextually appropriate, aligning with the structured schema of the task.

Furthermore, the simplicity of this approach—using similarity search and a straightforward prompt structure—allowed us to maintain efficiency while achieving high accuracy. Unlike more complex reasoning techniques such as Chain-of-Thought (CoT), our implementation focused on minimizing computational overhead without sacrificing precision. This was particularly important in the healthcare domain, where even minor inaccuracies could lead to significant risks. By combining semantic chunking with RAG, we ensured that the generated outputs were both accurate and contextually appropriate, aligning with the structured schema of the task.

Upon mounting the file system that serves as the source for our data, the following steps are undertaken to structure and process the information systematically:

1. Source File (PDF) Specification: Leveraging the PyPDFLoader class from the pypdf package within the Python programming language, specifically executed in a Jupyter notebook environment, we systematically manage the content of PDF files. This includes the extraction of pertinent information from a predefined directory housing a curated selection of clinical documents across various categories. The subsequent utilization of PyPDFLoader facilitates the streamlined processing of content from each document with the corresponding Language Model (LLM). The semantic chunker was employed to divide the extracted text into semantically coherent segments. This approach was critical in reducing irrelevant or fragmented information, ensuring only contextually relevant chunks were used in the retrieval and generation process.

2. Query Schema Creation: The formulation of a structured query is conducted to generate a JSON-style schema, systematically aligned with key patient attributes such as name, age, diagnostic tests, diagnosis, and medication. Ensuring adherence to this specified schema is imperative. Python's object-oriented programming paradigm, implemented within a Jupyter notebook, is instrumental in defining the class that underpins this schema, thereby ensuring seamless data extraction and subsequent processing.

3. Prompt Formatting: Prior to submitting the prompt for processing by the LLM, we rigorously format it to align precisely with the schema defined by Pydantic. This formatting process, executed in Python and complemented by the Pydantic library for data validation within the Jupyter notebook framework, ensures that the response from the LLM strictly adheres to the predefined schema.

4. Model Interaction: The transmission of the formatted prompt to the Language Model (LLM) is facilitated through serialization. This serialization process is executed using Python, either through the OpenAI API or the Ollama library, contingent upon the specific case. The LLM, embedded within a Jupyter notebook, retrieves embeddings and pertinent data, applying predefined processing rules from various data models. The culmination of this interaction is directed towards a .bin file, serving as a repository for valuable embeddings. The retrieved chunks were directly appended to the input prompt as contextual information for the LLM. This facilitated structured and accurate response generation aligned with the predefined JSON schema.

5. Result Formatting: The outcomes of the prompt, critically, are not processed as plain text but undergo transformation into JSON format. This strategic conversion enhances clarity and eases interpretation, ensuring a structured representation of the results. The Python-based implementation, within the Jupyter notebook environment, facilitates subsequent processing and detailed analysis.

The execution of this methodology was conducted on high-performance hardware equipped with advanced processors, sufficient memory, extensive storage, and specialized hardware optimized for accelerating Large Language Model (LLM) computations. The computational environment, seamlessly integrated with the efficiency of a Jupyter notebook, constitutes a critical component of our execution framework.

Our comprehensive workflow unfolds within the structured formalism of Python programming, harnessing the versatile capabilities of a Jupyter notebook. This robust combination facilitates not only the extraction and structuring of data from medical reports but also ensures dynamic and efficient handling throughout the entirety of the process. The integration of advanced hardware, including an Intel Core i7 processor, an NVIDIA GeForce RTX graphics card, and 16 GB of RAM, provided a solid computational foundation for executing the complex tasks involved.

A critical component of this workflow was the use of a semantic chunker, which ensured that the data segments processed and retrieved maintained semantic coherence. This step significantly improved the reliability of the retrieval and generation processes, particularly in a healthcare context where accuracy and contextual relevance are paramount. By prioritizing semantically meaningful chunks, the methodology reduced the risk of hallucinations and irrelevant outputs, thus aligning the generated results more closely with the intended objectives.

The decision to employ semantic chunking was informed by its demonstrated advantages in prior studies, which highlight its ability to enhance precision in information extraction and semantic analysis, reduce time and memory costs, and improve the handling of complex structures [31, 32]. These benefits align closely with the requirements of our task, where maintaining semantic coherence and contextual relevance is essential for ensuring the accuracy and reliability of the generated outputs.

For a comprehensive evaluation of the model's performance, including detailed metrics such as accuracy, precision, recall, and F1-score, please refer to the Evaluation and Results sections. These sections provide an in-depth analysis of how RAG improves the accuracy and reliability of the generated outputs, particularly in the context of healthcare applications where precision is paramount.

The evaluation of our RAG-based approach focused primarily on the generation component, as detailed in the Evaluation and Results sections. Metrics such as accuracy, precision, recall, and F1-score were used to assess the quality of the final outputs, ensuring that the generated responses were both accurate and contextually appropriate.

For the retrieval component, we employed a pragmatic approach to select the top 3 chunks ($k=3$) based on cosine similarity to the query. This decision was guided by the structure and size of the clinical documents, which typically consisted of approximately 2 pages with a consistent format. Given this limited scope, retrieving 3 chunks provided a sufficiently strict yet manageable amount of context for the generation process. This approach minimized the risk of including irrelevant or fragmented information in the prompt while ensuring that the most relevant content was prioritized.

While a separate evaluation of the retrieval process (e.g., using metrics like Recall@K [33] or Mean Reciprocal Rank [34]) was not conducted, the observed performance in the generation phase—coupled with the structured nature of the source documents—supports the effectiveness of our retrieval strategy. Future work could explore more granular evaluations of the retrieval component to further optimize the balance between chunk relevance and computational efficiency.

Given the specific nature of our task—extracting structured information from medical documents stored in external repositories—a comparison with a pure LLM prompting approach (without retrieval) is not applicable. Our methodology is designed to leverage the retrieval of relevant chunks from the documents themselves, ensuring that the generated outputs are grounded in the specific content of the source material. This approach is fundamentally different from traditional LLM prompting, which relies solely on the model's pre-trained knowledge and does not incorporate external document retrieval.

Moreover, fine-tuning the model with proprietary data was not considered necessary or viable for this study. Fine-tuning typically requires a large amount of annotated data, which can be costly and time-consuming to produce, especially in specialized domains like healthcare. Instead, our zero-shot prompting strategy, combined with semantic chunking and retrieval-augmented generation, provides a scalable and flexible solution for extracting structured information from medical documents without the need for extensive training data. This approach allows us to maintain high accuracy and reliability while minimizing computational overhead and resource requirements.

Code and Implementation Details

To ensure transparency and reproducibility, the implementation of this study, including preprocessing, semantic chunking, model prompting, and result generation, was conducted in Jupyter Notebooks using Python. The notebooks contain detailed steps for extracting structured information from unstructured medical reports and demonstrate the application of advanced language models (LLMs) in a clinical context. The complete codebase, including configuration parameters, prompt templates, and examples for executing retrieval-augmented generation (RAG) workflows, is publicly available on GitHub at <https://github.com/garciacarmonaam/llm-dataext-experiments>. This repository ensures that the methodology can be replicated or adapted to other datasets and scenarios.

Evaluation

To quantify the model's performance, we evaluated its outputs across the five categories of the JSON schematic framework (name, age, diagnosis, tests, and medications) using standard metrics derived from the confusion matrix. These metrics include accuracy, precision, recall, and the F1 score, which collectively provide a comprehensive assessment of the model's effectiveness in extracting structured information from medical reports.

The ground truth for evaluation was established by manually annotating a subset of the dataset, ensuring that each patient attribute (name, age, diagnosis, tests, and medications) was accurately labeled. This annotated dataset served as the reference for comparing the model's predictions. The distribution of entities in the ground truth varied across categories, with some categories (e.g., diagnoses and tests) having a higher frequency of positive instances compared to others (e.g., names and ages). This imbalance highlights the importance of using metrics like precision, recall, and F1 score, which are more informative than accuracy in scenarios with uneven class distributions.

Our evaluation approach aligns with the methodology used by Fornasiere et al., who employed Mistral 7B for medical information extraction tasks, including medication and timeline extraction. Similar to our study, they utilized standard metrics such as precision, recall, and F1-score to evaluate model performance. However, while our study focused on a zero-shot prompting approach, explored multiple prompting strategies, including zero-shot, few-shot, and sequential prompting. Their results demonstrated that few-shot and sequential prompting significantly improved model performance, particularly in tasks requiring detailed information extraction, such as identifying medication dosage and frequency.

In terms of performance, they reported an F1-score of 0.683 for medication extraction using a few-shot approach with JSON output, which is comparable to our model's performance in similar categories. However, their study also highlighted challenges in extracting full medication details, achieving lower F1-scores for tasks involving dosage and frequency extraction. This aligns with our findings, where the model struggled to achieve high recall in categories like names and ages, likely due to the variability and complexity of the data. [25]

While our study primarily employed a zero-shot approach, fine-tuning represents a powerful alternative for enhancing model performance in domain-specific tasks like medical information extraction. Fine-tuning involves adapting a pre-trained language model (LM) to a specific domain by continuing its training on a smaller, task-specific dataset. This process allows the model to better

capture domain-specific terminology, context, and nuances, which are critical in healthcare applications.

For example, models like BioBERT and ClinicalBERT have demonstrated the effectiveness of fine-tuning in medical NLP tasks. BioBERT, a domain-specific adaptation of BERT, was fine-tuned on biomedical text corpora and achieved state-of-the-art performance in tasks such as named entity recognition (NER) and relation extraction in the biomedical domain. Similarly, ClinicalBERT, fine-tuned on clinical notes from electronic health records (EHRs), has shown superior performance in extracting clinical concepts and predicting patient outcomes. These models highlight the strengths of fine-tuning, particularly its ability to improve precision and recall in complex, domain-specific tasks.

However, fine-tuning also has its limitations. It requires a substantial amount of annotated data, which can be costly and time-consuming to produce, especially in specialized domains like healthcare. Additionally, fine-tuned models may suffer from overfitting if the training dataset is too small or not representative of the broader domain. This can limit their generalizability to new or unseen data. Despite these challenges, fine-tuning remains a valuable approach for improving model performance in tasks where domain-specific knowledge is critical. [26]

In a recent study by Ntinopoulos et al. (2025), the performance of multiple large language models (LLMs) was evaluated for data extraction from unstructured and semi-structured electronic health records (EHRs). Their findings revealed that models like Claude 3.0 Opus, GPT-4, and Llama 3-70b achieved outstanding accuracy (>0.98) in both entity extraction and binary classification tasks. These results are consistent with our observations, where the model demonstrated high precision and recall in extracting structured information from unstructured PDFs. However, Ntinopoulos et al. also highlighted challenges in handling long, unstructured texts, particularly when relevant information is scattered throughout the document. This aligns with our findings, where the model struggled with categories like ages and medications, likely due to the variability and complexity of the data.

Specifically, the variability in how ages and medications are expressed (e.g., "45 años" vs. "45 y/o" or "Paracetamol" vs. "Acetaminofén"), combined with the lack of explicit contextual cues, makes these categories particularly challenging to extract accurately. Additionally, the dispersion of relevant information across the document further complicates the extraction process. In unstructured PDFs, critical details such as ages or medications may appear in different sections, often without clear labels or consistent formatting. This contrasts with more structured data, where information is typically organized in predictable ways (e.g., tables or labeled fields). The need for the model to navigate and interpret such dispersed information adds another layer of complexity.

Ntinopoulos et al. also emphasized the importance of response consistency across multiple iterations of the same prompt, a factor that we consider critical for ensuring the reliability of our model in real-world applications. In their study, models like Claude 3.0 Opus and GPT-4 demonstrated high consistency, with minimal variation in responses across multiple runs. This is particularly important in clinical settings, where inconsistent outputs could lead to errors in patient care or data analysis. While our current evaluation focuses on accuracy and recall, future work will include consistency assessments

to further validate the model's robustness. This aligns with the broader trend in the field, where consistency is increasingly recognized as a key metric for evaluating the reliability of LLMs in healthcare applications. [27]

To complement these consistency considerations, we employed standard evaluation metrics to quantify the model's performance. Accuracy measures the overall correctness of predictions, precision evaluates the relevance of the extracted data, recall assesses the system's ability to capture all pertinent information, and the F1 score provides a balanced measure that accounts for both precision and recall. These metrics are particularly useful for evaluating performance in scenarios with imbalanced data distributions, ensuring a robust assessment of the model's capabilities.

The JSON schema served as the foundation for structuring the extracted data, ensuring consistency and alignment with key patient attributes. By adhering to this schema, the model's outputs were systematically organized, facilitating both evaluation and integration into downstream applications. This structured approach not only streamlined the extraction process but also enabled a clear and standardized framework for assessing performance across diverse categories.

By employing these metrics and leveraging the JSON schema, our evaluation offers a detailed understanding of the model's performance, highlighting its strengths and areas for improvement in extracting and structuring data from medical reports. While fine-tuning presents a promising avenue for further performance gains, our zero-shot approach provides a scalable and flexible solution for medical information extraction, particularly in scenarios where annotated training data is limited.

Ethical Considerations

This study utilized anonymized clinical cases, ensuring that no personally identifiable information was included. As the dataset comprised fully de-identified cases prepared in accordance with institutional guidelines, no additional ethics review board approval was required. The anonymization process strictly followed established protocols to guarantee privacy and confidentiality, upholding the highest ethical standards for research involving secondary analysis of medical data.

Results

The evaluation of the models is presented in Table 2, which summarizes their performance across specific medical data categories: names (N), ages (A), diagnoses (D), tests (T), and medication (M). Key metrics such as Accuracy, Precision, Recall, and F1-score are provided, alongside an overall average (AVG) calculated across all categories.

As noted in the Evaluation section, the ground truth was established through manual annotation, and the distribution of entities varied significantly across categories. For instance, diagnoses and tests had a higher frequency of positive instances, while names and ages were less frequent. This imbalance underscores the importance of relying on metrics like precision, recall, and F1 score, which provide

a more nuanced understanding of model performance than accuracy alone.

Table 2. Correct scores per categories of different Large Language Models.

Model	Category	Accuracy	Precision	Recall	F1
GPT-4o	N	0.860	0.500	0.143	0.222
	A	0.970	1.000	0.970	0.985
	D	0.890	0.899	0.989	0.942
	T	0.950	0.949	1.000	0.974
	M	0.900	0.953	0.932	0.943
	AVG	0.914	0.860	0.807	0.813
LLaMa 3	N	0.370	0.088	0.313	0.137
	A	0.580	0.962	0.560	0.708
	D	0.750	0.888	0.816	0.850
	T	0.740	0.899	0.798	0.845
	M	0.700	0.983	0.667	0.795
	AVG	0.628	0.764	0.631	0.667
LLaMa 3.1	N	0.470	0.059	0.375	0.102
	A	0.510	0.940	0.505	0.657
	D	0.730	0.986	0.737	0.844
	T	0.710	0.947	0.740	0.830
	M	0.680	1.000	0.624	0.768
	AVG	0.620	0.786	0.596	0.640
Gemma	N	0.606	0.051	0.500	0.093
	A	0.485	0.957	0.473	0.633
	D	0.758	0.987	0.763	0.860
	T	0.707	0.985	0.702	0.820
	M	0.735	0.983	0.702	0.819
	AVG	0.658	0.793	0.628	0.645
Gemma 2	N	0.800	0.167	1.000	0.286
	A	0.710	0.945	0.734	0.826
	D	0.990	0.990	1.000	0.995
	T	0.980	1.000	0.980	0.990
	M	0.850	0.974	0.851	0.908
	AVG	0.800	0.167	1.000	0.286
Qwen 2	N	0.470	0.023	0.091	0.036
	A	0.400	0.925	0.394	0.552
	D	0.800	1.000	0.800	0.889
	T	0.800	0.988	0.806	0.888
	M	0.740	1.000	0.667	0.800
	AVG	0.642	0.787	0.551	0.633

Qwen 2.5	N	0.580	0.073	0.429	0.125
	A	0.560	0.847	0.588	0.694
	D	0.980	0.980	1.000	0.990
	T	0.930	0.979	0.949	0.964
	M	0.820	0.970	0.802	0.878
	AVG	0.774	0.770	0.754	0.730

An analysis of the results reveals significant variations in performance both between models and within individual categories. For instance, GPT-4o demonstrates outstanding overall performance, with an average accuracy of 0.914 and an F1-score of 0.813. However, it performs notably poorly in the names category, achieving an F1-score of 0.222, suggesting that the model struggles to process textual entities that are complex or inconsistent.

In contrast, Gemma 2 excels in categories such as diagnoses, achieving an F1-score of 0.995, and tests, with an F1-score of 0.990, showing high consistency in these critical areas. Nevertheless, its low performance in names, with a precision of 0.167, indicates a lack of balance across categories, which may limit its application in scenarios requiring the extraction of diverse types of sensitive data.

Models such as LLaMa 3 and LLaMa 3.1 exhibit similar patterns: relatively stable performance in diagnoses and tests but marked deficiencies in names and ages. For instance, LLaMa 3 achieves an F1-score of just 0.137 in names, while reaching an F1-score of 0.850 in diagnoses. This imbalance suggests that these architectures are less effective across all categories, potentially due to biases in training data or inherent limitations in their model design.

An interesting case is Qwen 2.5, which achieves a competitive average F1-score of 0.730 and strong performance in diagnoses, with an F1-score of 0.990, and tests, with an F1-score of 0.964. However, its performance in names (F1-score: 0.125) highlights a common trend across the evaluated models: significant challenges in this category, potentially due to the complexity and variability of names in medical contexts.

These results reflect the challenges posed by a zero-shot prompting approach, in which the models were tasked with extracting structured information without prior task-specific fine-tuning. While this method demonstrates the flexibility and adaptability of the models, it may also exacerbate limitations in categories requiring more nuanced understanding or specialized training, such as names and ages.

Overall, while the average metrics provide a general view of performance, the discrepancies across specific categories underscore the need for more specialized approaches to ensure consistent performance in medical applications. This analysis emphasizes the importance of optimizing both the models and prompting strategies to address the identified weaknesses and ensure reliability in real-world scenarios.

For the category of names, the datasets used in this study did not include actual personal identifiers due to anonymization. Instead, references to the absence of names (e.g., 'not available') or generic mentions of a person were included. The consistently poor performance of the models in this category indicates a limitation in

recognizing or interpreting such generic references within the text. This suggests that the models struggled with the ambiguity and variability introduced by the anonymized data.

Discussion

Among the evaluated models, GPT-4o demonstrated the highest overall performance, achieving an average score of 91.4% across all assessed categories. Each individual category score exceeded 80 points (out of one hundred), highlighting the model's consistency and robustness. In particular, GPT-4o excelled in accuracy, precision, and F1 score, especially in extracting age, diagnosis, and tests information. However, its recall, while satisfactory overall, was not as high as its other metrics, with Gemma 2 demonstrating superior recall rates in some categories.

A deeper analysis reveals that tasks like name extraction posed challenges for all models, particularly due to the anonymization of the dataset, which used placeholder or fictional names. This led to true negatives rather than errors, reflecting a limitation inherent to the dataset design rather than a failure of the models. Tasks like medication and diagnosis extraction, on the other hand, benefited from consistent terminologies and clearer patterns in the data, enabling models such as GPT-4o and Gemma 2 to achieve near-perfect precision and recall in these areas.

The anonymization of patient narratives in this study presented a unique challenge for the models in the names category. Rather than extracting explicit names, the models were tasked with identifying placeholders or generic references. This experimental setup, while necessary for data privacy, may not fully represent real-world scenarios where explicit personal identifiers are often present. Consequently, the results in this category should be interpreted with this limitation in mind.

From an ethical and legal perspective, privacy is a fundamental concern when handling medical data. Privacy can be interpreted as intrinsic to the right to property, which extends beyond tangible assets to include personal data such as health, economic, social, and nutritional information. Individuals often seek to control the extent to which external entities can access their personal data, particularly in sensitive domains like healthcare.

However, scientific and technological research faces a significant dilemma. While the right to privacy must be rigorously defended, the trial-and-error phases inherent in advancements in fields such as computer science, medicine, and pharmacology often require experiments with real-world data. This tension underscores the importance of strategies like anonymization and pseudonymization, which allow researchers to work with sensitive data while protecting individual identities.

Despite the benefits of open data for research and innovation, there is a lack of understanding about these strategies and the potential of open data to enhance scientific progress. Not all hospitals or institutions have open records suitable for experimentation, and ethical considerations often limit the availability of medical data for research purposes. In this study, the dataset of 100 Spanish medical reports was carefully anonymized to ensure compliance with ethical standards while enabling meaningful analysis. The decision to use anonymized data, rather than making the dataset publicly available, reflects the need to balance the advancement of medical research with the protection of patient privacy.

Gemma 2 followed closely behind GPT-4o, with an average score around 80%. Its performance was particularly notable in the diagnosis and tests categories, where it achieved recall rates of 1.000 and near-perfect accuracy. This suggests that Gemma 2 is well-suited for tasks requiring exhaustive retrieval of relevant information. However, its precision in the name category remained low, reflecting ongoing challenges in handling anonymized or placeholder data.

LLaMA 3 and LLaMA 3.1 presented intermediate performance, with average scores of 62.8% and 62%, respectively. Both models showed relative strengths in extracting diagnostic and test-related information, achieving moderate recall and F1 scores in these categories. However, their performance in extracting names and age data was weaker, likely due to variability in the data and limitations in their contextual understanding. The slight improvements in LLaMA 3.1 indicate potential benefits from iterative refinement in model architecture.

Qwen 2 and Qwen 2.5 demonstrated similar trends, with average scores of 64% and 77%, respectively. Qwen 2 excelled in tasks like diagnosis and tests, achieving perfect precision, but struggled significantly in name extraction due to placeholder data. Qwen 2.5 improved on these results, particularly in recall for diagnosis and tests, highlighting its potential for more complex retrieval tasks. Nevertheless, both models require further development to address challenges in handling diverse data categories effectively.

Meta models exhibited acceptable accuracy for categories like diagnosis, tests, and medications but faced difficulties in extracting names and age data. The significant class imbalance and variability in these categories adversely impacted their F1 scores. These results underscore the need for additional fine-tuning or hybrid approaches to enhance their performance in scenarios involving diverse and unstructured medical data.

The variability in performance across categories reflects the inherent challenges of applying LLMs to a domain as diverse as clinical medicine. The dataset utilized in this study spans multiple medical specialties and includes a wide range of clinical conditions, medications, and terminologies. While this diversity enhances the generalizability of the findings, it also introduces complexities that may not exist in more homogeneous datasets. For example, categories like medication and diagnosis benefit from the relative uniformity of medical terminology, while names and ages are inherently more variable due to anonymization and differences in reporting formats. Future efforts should explore the impact of dataset composition on model performance, particularly when applied to real-world clinical data.

From a broader perspective, these findings emphasize the adaptability of LLMs for extracting structured information from unstructured medical reports. Compared to traditional rule-based systems, LLMs provide greater flexibility and scalability, enabling them to handle a wide range of tasks and data formats. However, hybrid approaches that combine rule-based methods with generative capabilities could address some of the current limitations, particularly in high-stakes tasks like name extraction.

The use of semantic chunking and retrieval-augmented generation (RAG) in this study demonstrates the effectiveness of context-preserving techniques in minimizing hallucinations and improving result relevance. By integrating retrieved data directly into

the prompt, the models were able to generate structured outputs aligned with the predefined schema. This approach highlights the importance of carefully designed preprocessing steps to ensure consistent and reliable outputs.

The implications for clinical workflows are significant. By automating the extraction of critical patient information, LLMs reduce the cognitive load on healthcare professionals, streamline clinical workflows, and enable faster decision-making. These advantages are particularly evident in time-sensitive scenarios, where efficient information retrieval can make a substantial difference. However, the practical scalability of these solutions in resource-constrained environments remains an open question. Future work should investigate how these models can be adapted for deployment in settings with limited computational resources, ensuring their broader applicability and impact.

Despite these strengths, challenges remain. Future research should focus on addressing edge cases, such as ambiguous or inconsistent data, and on optimizing models for tasks requiring entity-specific recognition. Additionally, expanding error analysis to cover more granular categories and integrating domain-specific fine-tuning can further enhance the applicability of LLMs in healthcare settings.

This approach highlights the challenges of working with sensitive data in healthcare research and underscores the importance of developing robust frameworks for data anonymization and access control. Future work should focus on creating standardized protocols for data sharing that prioritize both innovation and ethical responsibility.

Data Availability

The data used in this study consist of anonymized clinical cases sourced from reputable organizations, including the Spanish Society of Internal Medicine (<https://www.fesemi.org>), the Asturian Society of Family and Community Medicine (<https://samfyc.org>), the Spanish Society of Cardiology (<https://secardiologia.es>), and the Faculty of Medicine at Francisco Marroquín University (<https://medicina.ufm.edu>). Given the sensitive nature of the data and their anonymization, they are not publicly available. However, further details about the data sources or methodologies used for anonymization can be provided upon reasonable request to the corresponding author.

Conclusions

This study explored various SCHEMA-TO-JSON strategies, leveraging the capabilities of large language models (LLMs) to organize and extract information from medical reports based on a JSON schema framework implemented using PyDantic. This approach aimed to systematically structure clinical data, transforming unstructured narratives into a standardized format. The methodology proved effective in organizing domain-specific healthcare information, laying a robust foundation for the development of tailored data models.

The experimental results demonstrate that the LLMs employed can effectively extract relevant information from medical histories. The high scores achieved in categories such as diagnoses and pharmacological data underscore the potential of these models to handle complex medical information. This aligns with findings in related studies, such as syntactic network analysis and knowledge

graph frameworks, confirming the utility of advanced natural language processing (NLP) techniques in the medical domain.

However, the study also highlights key areas for improvement:

- Challenges with Personal Details: OpenAI's models, despite their high overall performance, show inconsistencies in extracting specific details such as names and ages. These limitations are amplified in anonymized or pseudonymized contexts, where implicit or indirect references add complexity.
- Model Variability: Models such as Gemma 2 and Qwen 2.5 exhibit strong performance in diagnostic and pharmacological categories but share similar challenges in handling personal details. Meta's models require substantial improvement across multiple categories, suggesting a broader scope for refinement.

These findings emphasize the need for further optimization of LLMs in domain-specific applications, particularly when addressing sensitive or nuanced categories of data. Incorporating additional training focused on these challenges or integrating external knowledge sources, such as knowledge graphs, may enhance the precision and adaptability of these models.

This work underscores the importance of deploying advanced NLP strategies to improve information retrieval and analysis in the medical domain. By addressing the inherent challenges of structured and unstructured data, this study contributes to the ongoing development of models capable of navigating and interpreting complex clinical information more effectively. Future work will focus on refining the extraction of sensitive details and exploring the integration of complementary techniques to enhance the overall robustness and reliability of these systems.

Author Contributions: Conceptualization, A.C., M.P., E.P. and J.B.; methodology, A.C., M.P., and E.P.; software, A.C., M.P., and E.P.; validation, A.C., M.P., and E.P.; formal analysis, A.C., M.P., and E.P.; investigation, A.C., M.P., and E.P.; resources, E.P., and J.B.; data curation, A.C., M.P., and E.P.; writing—original draft preparation, A.C. and E.P.; writing—review and editing, A.C., M.P., E.P., and J.B.; visualization, A.C. and E.P.; supervision, E.P. and J.B.; project administration, A.C., E.P. and J.B.; All authors have read and agreed to the published version of the manuscript."

Conflicts of Interest: The authors declare no conflicts of interest."

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

REFERENCES

- [1] Sageena, G., Sharma, M., & Kapur, A. (2021). Evolution of smart healthcare: Telemedicine during COVID-19 pandemic. *Journal of the Institution of Engineers (India): Series B*, 102(6), 1319–1324. <https://doi.org/10.1007/s40031-021-00568-8>
- [2] Yellowlees, P. M., Chorba, K., Burke Parish, M., Wynn-Jones, H., & Nafiz, N. (2010). Telemedicine can make healthcare greener. *Telemedicine and e-Health*, 16(2), 229–232. <https://doi.org/10.1089/tmj.2009.0105>
- [3] Raoof, S. S., & Durai, M. A. S. (2022). A comprehensive review on smart healthcare: Applications, paradigms, and challenges with case studies. SCOPE, Vellore Institute of Technology, Katpadi, Vellore, India. <https://doi.org/10.1155/2022/4822235>

- [4] Fijačko, N.; Gosak, L.; Štiglic, G.; et al. Can ChatGPT pass the life support exams without entering the American heart association course? *Resuscitation* 2023, 185, 109732. DOI: 10.1016/j.resuscitation.2023.109732.
- [5] Karabacak, M.; Margetis, K. Embracing Large Language Models for Medical Applications: Opportunities and Challenges. *Cureus* 2023. DOI: 10.7759/cureus.39305.
- [6] Stade, E.; Stirman, S.; Ungar, L.; et al. Large language models could change the future of behavioral healthcare: a proposal for responsible development and evaluation. *npj Mental Health Research* 2024, 3, 12. DOI: 10.1038/s44184-024-00056-z.
- [7] Barnard, F.; Van Sittert, M.; Rambhatla, S. Self-Diagnosis and Large Language Models: A New Front for Medical Misinformation. *arXiv* 2023. Available online: <http://arxiv.org/abs/2307.04910> (accessed on Nov. 04, 2023).
- [8] Mirchandani, S.; Xia, F.; Florence, P.; et al. Large Language Models as General Pattern Machines. *arXiv* 2023. Available online: <http://arxiv.org/abs/2307.04721> (accessed on Nov. 04, 2023).
- [9] Shi, Y.; Wiggers, P.; Jonker, C. Towards recurrent neural networks language models with linguistic and contextual features. In *Proceedings of Interspeech 2012*, ISCA: Sep. 2012; pp. 1664–1667. DOI: 10.21437/Interspeech.2012-456.
- [10] Zhou, D.; Schärli, N.; Hou, L.; et al. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. *arXiv* 2023. Available online: <http://arxiv.org/abs/2205.10625> (accessed on Jul. 02, 2023).
- [11] Proceedings Volume 13404. (2024). Fifth International Conference on Control, Robotics, and Intelligent System (CCRIS 2024) (Paper 1340413). Accessed on <https://doi.org/10.1117/12.3049923>.
- [12] Stoffel, F.; Fischer, F. Using a knowledge graph data structure to analyze text documents (VAST challenge 2014 MC1). In *Proceedings of the 2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, IEEE: Paris: Oct. 2014; pp. 331–332. DOI: 10.1109/VAST.2014.7042551.
- [13] Žitnik, S.; Bajec, M. Text Mining in Medicine. In *Computational Medicine in Data Mining and Modeling*, Rakocevic, G.; Djukic, T.; Filipovic, N.; Milutinović, V., Eds.; Springer New York: New York, NY, USA, 2013; pp. 105–134. DOI: 10.1007/978-1-4614-8785-2_4.
- [14] Bisercic, A.; Nikolic, M.; van der Schaar, M.; et al. Interpretable Medical Diagnostics with Structured Data Extraction by Large Language Models. *arXiv* 2023. Available online: <http://arxiv.org/abs/2306.05052> (accessed on Nov. 05, 2023).
- [15] Bai, F., Kang, J., Stanovsky, G., Freitag, D., Dredze, M., & Ritter, A. (2024). Schema-driven information extraction from heterogeneous tables. In Findings of the Association for Computational Linguistics: EMNLP 2024 (pp. 10252–10273). Association for Computational Linguistics.
- [16] Yamashita, R.; Bird, K.; Yue-Cheng, C.; et al. Automated Identification and Measurement Extraction of Pancreatic Cystic Lesions from Free-Text Radiology Reports Using Natural Language Processing. *Radiol Artif Intell* 2022, 4(2), e210092. DOI: 10.1148/ryai.210092.
- [17] Yu, A.; Zhongyu, A.; Chloe, P.; et al. Automating Stroke Data Extraction From Free-Text Radiology Reports Using Natural Language Processing: Instrument Validation Study. *JMIR Med Inform* 2021, 9(5), e24381. DOI: 10.2196/24381.
- [18] Awal, R.; Zhang, L.; Agrawal, A. Investigating Prompting Techniques for Zero- and Few-Shot Visual Question Answering. *arXiv* 2023. Available online: <http://arxiv.org/abs/2306.09996> (accessed on Nov. 09, 2023).
- [19] Jiang, Z., Xu, F., Gao, L., Sun, Z., Liu, Q., Dwivedi-Yu, J., Yang, Y., Callan, J., & Neubig, G. (2023). Active retrieval augmented generation. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (pp. 7969–7992). Association for Computational Linguistics.
- [20] Ranjit, M., Ganapathy, G., Manuel, R., & Ganu, T. (2023). Retrieval augmented chest X-ray report generation using OpenAI GPT models. Proceedings of the 8th Machine Learning for Healthcare Conference, in Proceedings of Machine Learning Research, 219, 650–666. Accessed on <https://proceedings.mlr.press/v219/ranjit23a.html>.
- [21] Sivarakumar, S., Kelley, M., Samolyk-Mazzanti, A., Visweswaran, S., & Wang, Y. (2024). An empirical evaluation of prompting strategies for large language models in zero-shot clinical natural language processing: Algorithm development and validation study. *JMIR Medical Informatics*, 12, e55318. Accessed on <https://doi.org/10.2196/55318>.
- [22] Azamfirei, R., Kudchadkar, S. R., & Fackler, J. (2023). Large language models and the perils of their

- hallucinations. *Critical Care, 27*, 120. <https://doi.org/10.1186/s13054-023-04393-x>.
- [23] Shi, Y.; Wiggers, P.; Jonker, C. Towards recurrent neural networks language models with linguistic and contextual features. In *Proceedings of Interspeech 2012*, ISCA: Sep. 2012; pp. 1664–1667. DOI: 10.21437/Interspeech.2012-456.
- [24] Aggarwal, C.C.; Zhai, C. Information Extraction from Text. In *Mining Text Data*; Springer US: Boston, MA, 2012; pp. 11–41. DOI: 10.1007/978-1-4614-3223-4_2.
- [25] Raffaello Fornasiere, Nicolò Brunello, Vincenzo Scotti, and Mark Carman. 2024. Medical Information Extraction with Large Language Models. In *Proceedings of the 7th International Conference on Natural Language and Speech Processing (ICNLSP 2024)*, pages 456–466, Trento. Association for Computational Linguistics.
- [26] Nunes, M.; Bone, J.; Ferreira, J.; Elvas, L. Health Care Language Models and Their Fine-Tuning for Information Extraction: Scoping Review. *JMIR Med Inform* 2024;12:e60164. Available online: <https://medinform.jmir.org/2024/1/e60164>. DOI: 10.2196/60164
- [27] Ntinopoulos, V., Rodriguez Cetina Biefer, H., Tudorache, I., Papadopoulos, N., Odavic, D., Risteski, P., Haeussler, A., & Dzemali, O. (2025). Large language models for data extraction from unstructured and semi-structured electronic health records: a multiple model performance evaluation. *BMJ health & care informatics*, 32(1), e101139. Available online: <https://doi.org/10.1136/bmjhci-2024-101139>
- [28] Cooper, W. S., Gey, F. C., & Dabney, D. P. (1992). Probabilistic retrieval based on staged logistic regression. In *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval (SIGIR '92)* (pp. 198–210). Association for Computing Machinery. Available online: <https://doi.org/10.1145/133160.133199>
- [29] Abaho, M., & Alfaifi, Y. H. (2023). Select and Augment: Enhanced Dense Retrieval Knowledge Graph Augmentation. *Journal of Artificial Intelligence Research*, 78, 269–285. Accessed on <https://doi.org/10.1613/jair.1.14365>
- [30] Romanova, A. (2024). Enhancing NLP through GNN-driven knowledge graph rewiring and document classification. In *2024 35th Conference of Open Innovations Association (FRUCT)* (pp. 579–587). Tampere, Finland. Accessed on <https://doi.org/10.23919/FRUCT61870.2024.10516410>
- [31] Wang, X., Xu, Q., Ding, H., Wang, G., Yan, Y., & Chen, P. (2019). Control instruction information processing method based on semantic chunking.
- [32] Muszyńska, E., & Copestake, A. (2017). Realization of long sentences using chunking. *Proceedings of the 10th International Conference on Natural Language Generation* (pp. 218–222). Santiago de Compostela, Spain: Association for Computational Linguistics.
- [33] Patel, Y., Tolias, G., & Matas, J. (2022). Recall@ k surrogate loss with large batches and similarity mixup. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 7502–7511).
- [34] Chapelle, O., Metlzer, D., Zhang, Y., & Grinspan, P. (2009, November). Expected reciprocal rank for graded relevance. In *Proceedings of the 18th ACM conference on Information and knowledge management* (pp. 621–630).