



DALL-M: Context-aware clinical data augmentation with large language models

Chihcheng Hsieh^a, Catarina Moreira^{c,f,b}, Isabel Blanco Nobre^d, Sandra Costa Sousa^d,
Chun Ouyang^a, Margot Brereton^b, Joaquim Jorge^{e,f}, Jacinto C. Nascimento^e

^a School of Information Systems, Queensland University of Technology, 2 George Street, Brisbane, 4000, QLD, Australia

^b School of Computer Science, Queensland University of Technology, 2 George Street, Brisbane, 4000, QLD, Australia

^c Data Science Institute, University of Technology Sydney, 61 Broadway, Ultimo, Sydney, 2007, NSW, Australia

^d Imagiology Department, Grupo Lusitadas, Rua Abílio Mendes 12, Lisboa, 1500-458, Portugal

^e Instituto Superior Tecnico, Universidade de Lisboa, Av. Rovisco Pais 1, Lisboa, 1049-001, Portugal

^f INESC-ID, Av. Rovisco Pais 1, Lisboa, 1049-001, Portugal

ARTICLE INFO

Dataset link: <https://github.com/ChihchengHsieh/DALL-M>

Keywords:

Clinical data augmentation
Large language models
Human-centered AI

ABSTRACT

X-ray images are vital in medical diagnostics, but their effectiveness is limited without clinical context. Radiologists often find chest X-rays insufficient for diagnosing underlying diseases, necessitating the integration of structured clinical features with radiology reports.

To address this, we introduce DALL-M, a novel framework that enhances clinical datasets by generating contextual synthetic data. DALL-M augments structured patient data, including vital signs (e.g., heart rate, oxygen saturation), radiology findings (e.g., lesion presence), and demographic factors. It integrates this tabular data with contextual knowledge extracted from radiology reports and domain-specific resources (e.g., Radiopaedia, Wikipedia), ensuring clinical consistency and reliability.

DALL-M follows a three-phase process: (i) clinical context storage, (ii) expert query generation, and (iii) context-aware feature augmentation. Using large language models (LLMs), it generates both contextual synthetic values for existing clinical features and entirely new, clinically relevant features.

Applied to 799 cases from the MIMIC-IV dataset, DALL-M expanded the original 9 clinical features to 91. Empirical validation with machine learning models – including Decision Trees, Random Forests, XGBoost, and TabNET – demonstrated a 16.5% improvement in F1 score and a 25% increase in Precision and Recall.

DALL-M bridges an important gap in clinical data augmentation by preserving data integrity while enhancing predictive modeling in healthcare. Our results show that integrating LLM-generated synthetic features significantly improves model performance, making DALL-M a scalable and practical approach for AI-driven medical diagnostics.

1. Introduction

Data augmentation is a fundamental technique in deep learning designed to enhance the diversity and robustness of training datasets. In image processing, techniques such as translation, rotation, scaling, and flipping introduce new patterns and variations, helping models generalize to unseen data. Similarly, in natural language processing (NLP), methods like word masking, substitution, and reordering improve a model's ability to capture contextual relationships. However, augmenting structured tabular data in clinical settings – such as patient vitals, laboratory values, and imaging-derived features – presents unique challenges [1,2]. Unlike conventional augmentation techniques, medical

data requires preserving the relationships between features (e.g., ensuring consistency between oxygen saturation and respiratory rate) to maintain clinical validity. Our approach ensures these relationships remain intact by integrating structured patient data with contextual information from radiology reports and trusted clinical knowledge bases.

Clinical tabular data is highly structured and context-sensitive. Each data point, such as patient demographics, medical conditions, or test results, forms part of a more extensive, interdependent system. Modifying these features without preserving their inherent relationships can introduce inconsistencies that compromise clinical validity. [3,4]. For instance, if a patient's blood pressure is altered without adjusting

* Corresponding author at: Data Science Institute, University of Technology Sydney, 61 Broadway, Ultimo, Sydney, 2007, NSW, Australia.
E-mail address: catarina.pintomoreira@uts.edu.au (C. Moreira).

related factors such as age or comorbidities, the dataset may no longer align with real-world medical patterns. Such discrepancies can mislead machine learning models, resulting in incorrect risk assessments or inaccurate diagnostic predictions [5,6].

Unlike traditional augmentation methods that introduce randomness, our proposed framework ensures that newly generated data maintains medical integrity, preserving the logical and clinical consistency of real patient data. This is achieved through domain-aware augmentation techniques that incorporate structured medical knowledge.

1.1. Research gap

Traditional augmentation methods, such as random shuffling or noise injection, often disrupt the complex relationships between clinical features, undermining the predictive power of machine learning models. This is especially problematic in healthcare, where patient safety and diagnostic accuracy are paramount. Introducing noise into clinical metrics – like blood pressure or glucose levels – may simulate pathological conditions that do not exist, resulting in misleading patterns during model training. Such distortions pose significant risks, potentially leading to diagnostic errors or inappropriate treatment recommendations, directly affecting patient outcomes.

Moreover, conventional augmentation methods fail to capture the interdependencies between medical features. For instance, artificially increasing a patient’s heart rate without adjusting oxygen saturation or respiratory rate can introduce clinically implausible cases. Similarly, duplicating data through synthetic oversampling does not add new knowledge, whereas DALL-M derives medically valid features from clinical context using LLMs and retrieval-augmented knowledge.

Despite the availability of large medical datasets such as the UK Biobank [7], clinical practice often involves working with smaller, less diverse datasets. The high cost of collecting and curating medical data exacerbates this issue, limiting many institutions’ ability to develop robust deep learning models. Deep learning thrives on vast amounts of diverse data, and when clinical datasets are limited in size or variety, the ability to train models effectively is hindered. These limitations highlight a critical gap in how we approach data augmentation for clinical tabular data.

There is an urgent need for novel data augmentation techniques that can preserve the integrity and context of clinical datasets while enabling advanced machine learning models, even in data-constrained environments.

1.2. Proposed approach

To address these challenges, we propose a novel data augmentation framework that leverages the reasoning capabilities of Large Language Models (LLMs). While LLMs are traditionally applied in natural language tasks, their ability to infer patterns across structured, heterogeneous data makes them promising candidates for clinical data augmentation. LLMs can be fine-tuned to understand domain-specific knowledge, enabling them to generate new data points that are logically consistent with existing clinical information.

In this context, we define “inferable” features as those that maintain the relationships between clinical variables, adhering to established medical knowledge and logic. For example, an LLM trained on clinical data can infer additional symptoms or disease progressions that align with a patient’s existing conditions, generating new features that respect clinical constraints. However, fine-tuning LLMs for clinical tasks is challenging: the computational demands, potential biases inherited from pre-trained models, and the costs of frequent updates present significant obstacles.

Our approach mitigates these concerns by anchoring the LLM outputs to trusted clinical sources (such as Radiopedia) and rigorously validating the generated data through a classification task and expert reviews. This ensures that the augmented data is clinically valid and useful for downstream machine learning applications.

1.3. Contributions

This paper addresses the key research question: *How can Large Language Models effectively augment clinical datasets with inferable, contextually relevant features while preserving clinical integrity?* We answer this by introducing a novel framework for structured data augmentation, named DALL-M, which generates clinically coherent features through a three-phase process: clinical context storage, expert query generation, and context-aware feature augmentation.

Our contributions are as follows:

- 1. Novel framework for contextually relevant augmentation (DALL-M):** We introduce a framework that leverages LLMs to generate new clinical features, overcoming the limitations of traditional methods such as random shuffling or noise injection. Unlike these conventional methods, DALL-M ensures that the newly generated features are logically consistent with existing clinical data, maintaining data integrity.
- 2. Empirical validation of LLM-generated values for clinical features:** We present empirical evidence that DALL-M can generate clinically relevant and contextually accurate values for existing clinical features using LLMs. This validation shows that LLMs can synthesize patient-specific values that are statistically similar to real clinical data, thereby preserving the integrity and trustworthiness of the augmented dataset. By comparing LLM-generated values to ground-truth clinical measurements, we demonstrate that DALL-M outperforms traditional augmentation methods, which often introduce unrealistic or clinically inconsistent patterns.
- 3. Integration of LLM-augmented datasets with machine learning models:** We demonstrate how DALL-M enhances the performance of various machine learning models, including Decision Trees, Random Forests, XGBoost, and TabNET. By showing how contextually augmented features improve model robustness and accuracy, we address the challenge of applying deep learning to limited clinical data.
- 4. Anchoring to trusted clinical sources:** To mitigate the risk of hallucinations – where models generate fabricated or inaccurate clinical data – we anchor the LLM to trusted sources like Radiopaedia. This ensures the generated features are grounded in verified medical knowledge, safeguarding the clinical validity of the augmented data.

Our approach enhances model performance while preserving patient-centric data relationships by ensuring that the LLM-augmented clinical features are inferable, contextually relevant, and grounded in clinical logic. This method bridges the gap between limited clinical data and the need for robust machine learning models, offering a reliable, scalable solution for clinical data augmentation.

2. Related work

The augmentation of tabular data plays a crucial role in improving the robustness and generalizability of machine learning models, particularly in clinical applications where data availability is often limited. Traditional data augmentation methods for structured datasets have primarily relied on retrieval-based or generation-based approaches, each with distinct advantages and limitations. While retrieval-based techniques enrich datasets by incorporating external information, they are constrained by data availability and scalability issues. Conversely, generation-based methods leverage probabilistic models and deep learning to synthesize new data points but often struggle with overfitting and maintaining data integrity. In this section, we review conventional tabular data augmentation techniques, highlight their limitations, and position our proposed framework, DALL-M, within the broader landscape of machine learning-driven clinical data augmentation.

2.1. Conventional approaches to tabular data augmentation (TDA)

Conventional tabular data augmentation (TDA) can be roughly classified into retrieval-based or generation-based approaches [8]. The first category enhances the original table with realistic data sourced from table pools, while the second generates synthetic data without relying on external sources. Generation-based methods leverage generative models such as BERT [9], T5 [10], large language models like ChatGPT [11] and LLaMA [12], as well as deep learning architectures including variational autoencoders (VAE) [13], GANs [14], and diffusion models [15].

Retrieval-based approaches can be further divided into four main subcategories. The first is entity-based augmentation at the row level [16,17], which includes statistical [16,18,19], knowledge base-based [20–24], graph-based [25–27], and pre-trained language model (PLM)-based [28,29] techniques. The second subcategory, schema augmentation at the column level, involves value-based [30,31], semantic-based [32–36], and structure-based [27,37,38] joins. The third subcategory, cell completion at the cell level, uses attribute name [22,25,39], entity ID or name [22], and cell value [22,25,28,40,41] completion methods. Finally, table integration at the table level consists of compositional [22,25,39] and direct [42,43] integration strategies.

Generation-based approaches are categorized into four sub-classes. Record generation techniques can be either distribution-preserving [44–49] or designed for class-imbalance correction [50–53]. Feature construction techniques are further classified into explicit [54–56] and implicit [57,58] approaches. Cell imputation strategies involve statistical methods [59,60] or deep learning-based methods such as VAEs [13], GANs [14], and diffusion models [15]. Lastly, table synthesis aims to generate entire synthetic tables [61].

Despite the advancements in TDA, limitations persist. Retrieval-based methods extract related tables from table pools to perform data augmentation, enriching the original table with real external data. However, this approach depends on the availability of relevant information and requires preprocessing and indexing a potentially vast number of tables, raising scalability concerns. Additionally, the scarcity of labeled data in large-scale table pools suggests self-supervised approaches as a promising future direction. Furthermore, retrieval-based methods struggle with robust generalization. On the other hand, generation-based TDA synthesizes artificial data specifically for augmentation. Unlike retrieval-based strategies, these approaches do not require external data sources, eliminating preprocessing and indexing overhead. However, generation-based methods are often over-parameterized, leading to overfitting when working with small tables. Additionally, they lack interpretability and may produce hallucinated outputs, reducing the reliability of the generated data.

Our approach falls within the generation-based category, sharing similar concerns. However, the novelty of DALL-M lies in its context-aware feature generation, which leverages retrieval-augmented generation (RAG) techniques to enhance the clinical validity of synthetic data. Unlike conventional generation-based approaches that rely solely on statistical distributions or deep generative models, DALL-M incorporates structured domain knowledge from trusted clinical sources such as Radiopaedia and Wikipedia to inform and constrain the augmentation process.

DALL-M explicitly models the dependencies between clinical features, ensuring that generated synthetic data maintains medical consistency, such as heart rate variations remaining aligned with oxygen saturation and respiratory rate. This distinguishes our framework from traditional generative models, which often introduce inconsistencies by treating features independently. Another key advancement is DALL-M's ability to generate entirely new, clinically relevant features guided by expert-informed query generation. Instead of augmenting datasets with random perturbations, our method synthesizes additional patient attributes, enriching the dataset while preserving its interpretability. This mitigates the hallucination risk associated with large language models

by anchoring feature generation to validated medical knowledge.

Additionally, DALL-M is specifically designed for tabular data in clinical settings, where structured dependencies must be maintained. Many existing generative approaches focus on image and text data, making direct adaptation to clinical tabular datasets challenging. By integrating domain-specific retrieval, structured query generation, and feature augmentation, our method ensures that synthetic data is both realistic and clinically meaningful.

2.2. LLM-based approaches for data augmentation

Medical Large Language Models (Med-LLMs) are increasingly used in clinical workflows for tasks such as improving medical knowledge understanding, enhancing diagnostic accuracy, and providing personalized treatment recommendations. These capabilities contribute to precise decision-making, better patient care, and improved treatment outcomes. Several Med-LLMs, including ChiMed-GPT [62], Medical-GPT [63], HuatuoGPT-II [64], and ChatMed [65], have gained significant attention in biomedical research. Existing studies can be categorized into two main areas [1]: current Med-LLMs in the medical field and their role in supporting clinical and patient treatment.

Several existing Med-LLM methodologies have emerged. HuatuoGPT [64] utilizes reinforcement learning to align LLMs with real-world data. ClinicalT5 [66] adapts general LLMs for clinical text by pre-training on MIMIC-III data. ClinicalGPT [67] integrates medical records, domain-specific knowledge, and multi-round consultations. ChiMed-GPT [62] trains on domain-specific knowledge, while BioGPT [68] enhances biomedical text generation capabilities. Other models such as PubMedBERT [40], GatorTron [69], and Med-PaLM [70] focus on various aspects of clinical data processing, including electronic health records and multimodal medical analysis.

Despite the broad applications of Med-LLMs, this study proposes utilizing them in a novel way. Previous studies such as GReAT [2], HALO [3], and TEMED-LLM [4] have explored Med-LLMs for tabular data, but none have specifically tackled clinical data augmentation. DALL-M fills this gap by addressing how to generate realistic clinical features and infer new attributes. Our methodology integrates domain experts in evaluating generated features, ensuring that synthetic clinical features are meaningful and valuable for machine learning applications in medical diagnostics.

Table 1 provides a comparative overview of existing LLM-based approaches for medical data augmentation, highlighting their key focus areas, limitations, and how DALL-M addresses these gaps by generating new clinically relevant features and enhancing tabular data integration.

In summary, DALL-M pioneers LLM-driven tabular data augmentation by combining retrieval-augmented contextualization, structured knowledge integration, and expert-guided feature generation—capabilities absent in prior methods. Unlike traditional augmentation techniques that rely on random permutations or noise injection, DALL-M ensures clinical validity by dynamically retrieving and synthesizing multi-source medical knowledge (e.g., Radiopaedia, Wikipedia). It generates entirely new, context-aware clinical features, bridging the gap between text-based LLMs and structured data augmentation. By incorporating expert-driven constraints, DALL-M mitigates biases, hallucinations, and inconsistencies, leading to more interpretable, reliable, and clinically relevant synthetic data, ultimately improving AI-driven healthcare predictions and decision-making.

3. DALL-M: Data augmentation with LLMs

The proposed framework, DALL-M, aims to enhance clinical datasets by generating new, contextually relevant features using LLMs. The framework consists of three key stages: (1) Clinical Context Extraction and Storage, (2) Expert Input Queries and Prompt Generation, and (3) Context-Aware Feature Augmentation. Together, these stages form a

Table 1
Summary of LLM-based approaches for data augmentation.

Approach	Key Focus	Limitations	Role of DALL-M
ChiMed-GPT [62], MedicalGPT [63], HuatuogPT-II [64], ChatMed [65]	Adapt general-domain LLMs for clinical tasks	Primarily focused on text-based reasoning (e.g., question answer tasks)	Leverages LLMs in a novel way to enrich <i>tabular</i> data with new features and synthetic values
	Harness large-scale medical text corpora	Less attention to generating entirely new clinical features	Goes beyond question-answering by creating entirely new features relevant to specific radiology contexts
	Provide domain-specific responses (e.g., diagnoses)	Limited exploration of tabular data	Focuses on text and tabular data
ClinicalT5 [66]	Fine-tuned T5 for clinical text tasks (e.g., summarizing patient records)	Focus on generating or refining text-based outputs, not specialized in augmenting structured (tabular) clinical data with new features	DALL-M introduces a three-phase pipeline that leverages multiple data sources to create augmented tabular data
	Shows improved performance in typical NLP tasks such as entity recognition, summarization	Does not explore advanced context integration (e.g., from multiple sources like Radiopaedia, Wikipedia)	Adds clinically relevant features rather than just refining text
ClinicalGPT [67]	Integrates various info sources (medical records, domain knowledge, dialogues)	Emphasis on multi-round consultations and EHR text retrieval	DALL-M not only consolidates domain knowledge from multiple sources but also generates new, context-aware <i>tabular</i> features
	Enables context-rich responses in interactive clinical settings	No explicit mechanism for structured tabular feature augmentation or generation of new features	Uses a retrieval-augmented pipeline to ensure alignment of new features with established clinical knowledge
BioGPT [68] PubMedBERT [40] GatorTron [69]	Domain-trained LLMs	Specialized for textual data (biomedical corpora, EHR text) with minimal focus on synthetic feature generation in numerical or tabular formats	DALL-M embeds new numeric or categorical features into existing tabular records, preserving consistency and context
	Enhanced accuracy in medical text classification, knowledge extraction, etc	Do not address missing data or new feature creation in a robust data-augmentation setting	Addresses missing data scenarios and introduces novel features (e.g., additional vital stats, symptom descriptors)
Med-PaLM [70] MedAlpaca [41] LLaVA-Med [71]	Employ large-scale domain adaptation of general LLMs	Focus primarily on question-answer or generative tasks (report summarization, explanations)	DALL-M explicitly tackles the augmentation of <i>tabular</i> clinical data with entirely new features (not just text or imaging commentary)
	Incorporate specialized medical data Provide multi-modal understanding (images + text)	Limited exploration of how to augment <i>structured, tabular</i> clinical data with newly inferred features	Uses multi-modal and multi-source knowledge (X-rays + textual reports + Radiopaedia + Wikipedia) for robust feature generation
TDA (Noise Injections, Data Permut.)	Simple to implement for numeric data	Often breaks clinical consistency (random noise can create medically implausible combinations)	DALL-M preserves clinical logic by integrating domain knowledge from multiple sources and structured expert queries
	Can increase dataset size quickly for ML model training	Does not leverage domain knowledge to ensure realism or generate completely new features	-Ensures newly generated or augmented features remain medically plausible
DALL-M (Our approach)	Generates synthetic, context-relevant numeric values for missing data	Requires advanced retrieval-augmented methods and curated knowledge sources	Fills the gap of systematically augmenting tabular data with <i>both</i> realistic values for existing columns and <i>entirely new</i> clinically relevant features
	Creates entirely new features based on domain knowledge and retrieval augmented LLM queries	Dependent on LLM frameworks (e.g., GPT-4) for best performance, which can be computationally expensive	Provides robust, context-aware augmentation that improves downstream ML performance and addresses missing data issues
	Integrates multiple data sources (e.g., Radiopaedia, Wikipedia) and expert inputs	Ethical and privacy guidelines for synthetic data must be carefully observed	
	Demonstrates efficacy across multiple ML models (e.g., XGBoost, TabNet)		

comprehensive data augmentation workflow, ranging from the extraction of clinical context to the generation of enhanced features, all while leveraging the analytical capabilities of LLMs. These stages are illustrated in Fig. 1. The following sections provide detailed descriptions of each phase. To improve clarity, we introduce a running example throughout this section:

Consider a patient whose chest X-ray shows evidence of pleural effusion. The initial dataset only contains three clinical features for this patient: age (65 years old), temperature (38.2 °C), and oxygen saturation (92%). Our goal is to augment this patient's clinical data by generating additional relevant features while preserving medical integrity. DALL-M will first extract medical knowledge about pleural effusion (Phase I), use expert-informed queries to retrieve textual contents related to the patient's clinical context (Phase II), and

finally generate new synthetic and clinically meaningful features (Phase III).

3.1. Phase I: Clinical context extraction and storage

The primary objective of this phase is the extraction and structured storage of clinically relevant context and patient-specific medical history. Beginning with a particular patient case, the pipeline systematically builds a comprehensive knowledge repository that encapsulates detailed clinical information about the patient's conditions. In this context, "conditions" refer specifically to lesions identified via ground truth labels from the REFLACX dataset [72], a well-established radiological dataset that provides expert-annotated lesion labels as a foundation for subsequent knowledge extraction. A lesion in this

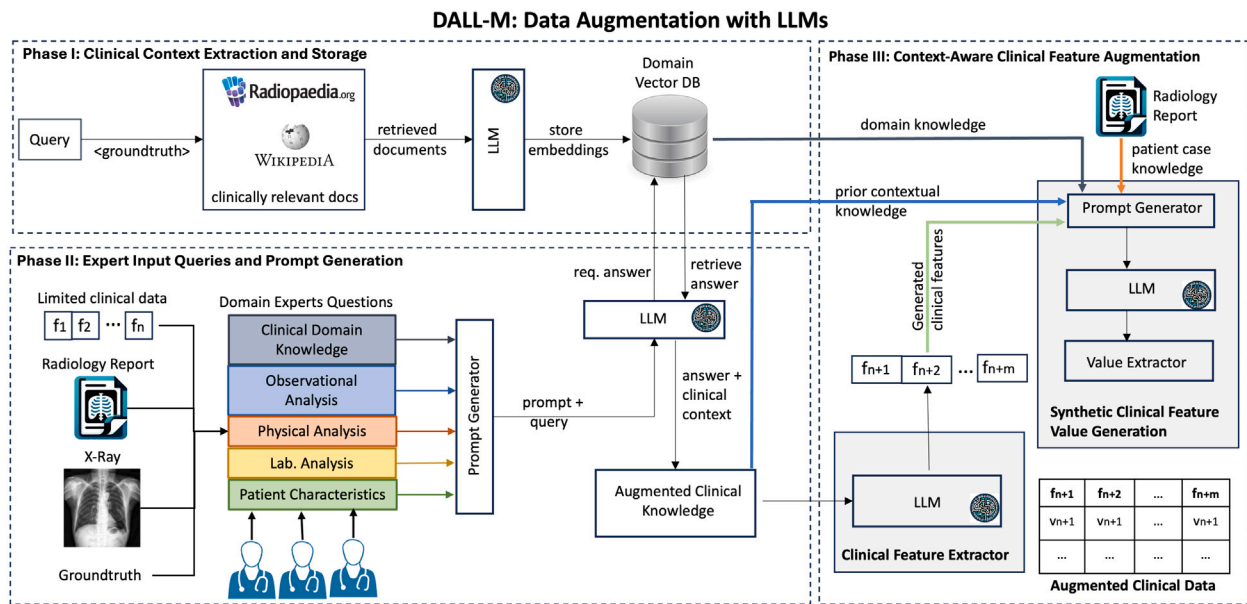


Fig. 1. Overview of the DALL-M framework for generating clinically relevant features using Large Language Models. The process consists of three main phases: (I) *Clinical Context Extraction and Storage*, where patient-specific contexts and clinical relationships are captured and stored; (II) *Expert Input Queries and Prompt Generation*, where medical experts provide contextual queries to guide the LLM; and (III) *Context-Aware Clinical Feature Augmentation*, where new features are generated in alignment with the clinical context, ensuring that the augmented data remains consistent and clinically valid.

context is an abnormality found in the X-ray, for instance, a lung nodule.

The process begins by leveraging these lesion labels to query two authoritative and distinct online resources: Radiopaedia [73] and Wikipedia [74]. Radiopaedia, a leading online platform dedicated to radiology, provides a robust collection of peer-reviewed radiological cases, detailed articles, and reference materials, all aimed at enhancing clinical knowledge and decision-making in radiology. Wikipedia, on the other hand, is a vast, freely accessible encyclopedia that covers a wide range of topics, offering a more general contextual understanding that can complement specialized medical knowledge. Both resources are queried to establish a broad and detailed clinical foundation for each patient case, ensuring coverage of both domain-specific and general knowledge relevant to the identified lesions.

We employ the Retrieval Augmented Generation (RAG) framework to systematically process and organize this information. RAG is a hybrid methodology that integrates traditional information retrieval techniques with the generative power of LLMs. The framework comprises two components: a retriever model and a generator model. The retriever model is responsible for identifying and extracting relevant documents from external sources (in this case, Radiopaedia and Wikipedia) based on the lesion labels. These documents are then passed to the generator model, which processes the retrieved information and generates semantically coherent, contextually enriched text that aligns with the clinical requirements of the task. The integration of real-time information retrieval with LLMs ensures that the generated responses are not only accurate but also context-aware, providing a comprehensive clinical picture tailored to the patient's case.

The output of the RAG process is then stored in a structured format to enable efficient querying and retrieval. After evaluating multiple database architectures, including relational and graph-based models, we determined that Neo4j (specifically utilizing its vector index capabilities) provides optimal performance for this task. Neo4j, a native graph database, is well-suited for capturing complex relationships between clinical entities, such as lesions, symptoms, and patient characteristics. The addition of vector indexing enhances this by enabling efficient similarity-based retrieval, ensuring semantically related information is easily accessible for subsequent phases. This configuration was meticulously applied to each of the five lesion labels extracted

from the REFLACX dataset, yielding a domain-specific clinical knowledge graph that stores relevant clinical context and ensures rapid and accurate retrieval for downstream analysis.

One of the key advantages of using a RAG approach is that it anchors LLM responses to authoritative medical sources (e.g., Radiopaedia, Wikipedia). Unlike free-form LLM text generation, this grounding mechanism significantly reduces hallucinations, as all generated features and values are derived from verified medical knowledge. This ensures that synthetic data remains clinically coherent and minimally prone to incorrect or misleading outputs.

Running Example (Phase I): The patient's X-ray is labeled with a lesion "pleural effusion". Our system retrieves detailed medical descriptions, causes, symptoms, and diagnostic indicators of pleural effusion from Radiopaedia and Wikipedia. This information is stored in a structured Neo4j database for later use. Key findings include: (1) Common causes such as heart failure and infection, (2) Typical symptoms like shortness of breath and chest pain, and (3) Diagnostic indicators such as blunting of the costophrenic angle.

To summarize, Phase I takes a specific clinical query (e.g., "pleural effusion"), retrieves the primary documents related to that query from Radiopaedia and Wikipedia and stores them in a database. This structured storage provides a strong foundation for clinical knowledge that will be used in later stages for data augmentation and feature generation.

3.2. Phase II: Expert input queries and prompt generation

After having the database populated with documents related to the user query, Phase II generates clinically relevant, contextually rich domain knowledge that mirrors radiologists' diagnostic process when assessing X-rays. To ensure medical validity, we consulted two board-certified radiologists, each with 8 to 15 years of experience in diagnostic imaging and chest X-ray interpretation. We conducted structured interviews with expert radiologists to understand the critical clinical questions that arise during their diagnostic workflow when they are presented with an X-ray. They aimed to distill their expertise into focused queries guiding subsequent data augmentation and feature generation.

Expert Queries

Through these interviews, we transcribed and analyzed expert responses using thematic coding to extract common diagnostic reasoning patterns. This enabled us to systematically derive a structured query framework, ensuring that each lesion type is assessed through a standardized set of clinically relevant questions. These questions form a structured framework to guide data retrieval and augmentation, ensuring the clinical insights generated are comprehensive and aligned with radiological best practices.

To automate the application of these queries, we implemented a dynamic template-based approach where lesion-specific terms (e.g., “atelectasis” or “pleural effusion”) are automatically inserted into predefined question formats. This ensures consistency and scalability in applying the expert queries across different patient cases.

General Clinical Domain Knowledge:

- What are the most common symptoms associated with {lesion}?
- What are the potential causes of {lesion}?

These questions help establish a broad understanding of the lesion, providing a general clinical overview that can be applied across various cases. A lesion in this context is an abnormality found in the X-ray, for instance, a lung nodule.

Observational Analysis:

- What are the patient’s symptoms that are relevant for {lesion}?

This ensures that the patient’s subjective experience is integrated into the clinical assessment, enriching the context of the diagnosis.

Physical Examination and Imaging:

- What are the relevant clinical signs for the etiological diagnosis of {lesion}?
- What are the relevant clinical characteristics for the etiological diagnosis of {lesion}?

These queries capture the objective clinical observations and imaging findings that radiologists rely on to characterize and diagnose a lesion.

Laboratory Data:

- What are the relevant laboratory data for the etiological diagnosis of {lesion}?

These questions allow for integrating laboratory results and additional imaging tests into the diagnostic process.

Patient Characteristics:

- What is the patient’s personal, relevant history for the etiological diagnosis of {lesion}?

These questions contextualize the clinical findings within the patient’s broader medical history and demographic profile.

These structured queries serve as inputs to our Retrieval-Augmented Generation framework, which retrieves relevant clinical information from Radiopaedia and Wikipedia. This approach ensures that the augmentation process is informed by domain knowledge, dynamically adapting to different lesions while maintaining consistency with radiological best practices.

Prompt Generation and Clinical Information Retrieval

Given these structured questions, the next step is to construct prompts that effectively query the contextual clinical database built in Phase I. Each question is translated into a prompt designed to elicit detailed responses from the LLMs through the RAG approach.

To achieve this, we apply a RAG process in the same two-step approach outlined in Phase I, wherein the retriever model identifies the most relevant documents from the domain-specific vector database, and

the generator model synthesizes contextually enriched responses from the retrieved material. This approach ensures the generated outputs are grounded in authoritative clinical knowledge, tailored to the patient’s specific case, and aligned with expert radiological practices.

The output of this phase is an Augmented Clinical Knowledge (ACK) corpus, a curated collection of clinical insights that corresponds to the answers to the expert queries. The ACK captures information about the patient’s condition, including detailed explanations of the lesions, potential symptoms, and patient-specific factors (such as comorbidities or risk factors like obesity) that could influence the presentation of the lesion. This corpus enriches the clinical context for each patient case and provides critical input for the next phase, where new clinical features will be generated based on this augmented knowledge. This ensures that the subsequent data augmentation and feature generation processes are clinically meaningful and grounded in real-world radiological expertise.

Running Example (Phase II): For pleural effusion, the expert queries retrieve information such as: (1) Symptoms: Shortness of breath, chest pain, cough; (2) Possible causes: Heart failure, pneumonia, malignancy; (3) Relevant clinical features: Decreased breath sounds, dullness to percussion, jugular venous distension. These findings are stored in a structured format and fed into the next phase.

3.3. Phase III: Context-aware feature augmentation

The primary objectives of this phase are twofold: (1) to identify and generate new clinically relevant features and (2) to assign meaningful values to these newly identified features. This process builds on the Augmented Clinical Knowledge corpus developed in Phase II, extracting novel clinical features that are contextually aligned with the patient’s condition. Extracting such features from unstructured textual data presents challenges, even with advanced techniques like named entity recognition, as medical text often includes nuanced and domain-specific terminology that is difficult to capture in its entirety.

To address these challenges, we leverage LLMs equipped with few-shot learning capabilities, allowing them to identify relevant features even with minimal labeled examples. The LLM can discern n distinct features by examining the documents generated in Phase II, each feature representing a clinically important aspect of the patient’s condition. These features are not merely repetitions of existing variables but represent new data dimensions that provide deeper insights into the patient’s health status.

However, the mere identification of relevant clinical features is insufficient for a comprehensive augmentation of the clinical dataset. Each feature must be populated with corresponding values to ensure its practical utility. To achieve this, we revisit the RAG framework, crafting prompts that utilize multiple critical sources of information:

1. **Prior contextual knowledge** derived from the ACK corpus in Phase II provides a comprehensive understanding of the patient’s condition.
2. **Newly identified clinical features**, ensuring that the values generated directly relate to the novel attributes extracted in this phase.
3. **The domain knowledge database established in Phase I** serves as a reference for validated medical information related to the lesion and patient.
4. **The patient’s radiology report and demographic details** offer personalized data for generating patient-specific feature values.

By integrating these four sources of information, we generate precise and contextually relevant values for the newly identified features. For example, if a new feature related to lesion morphology is identified,

values would be generated based on both the clinical descriptions in the radiology report and the broader domain knowledge available in the ACK corpus.

When processed by the LLM within the RAG framework, these prompts return a comprehensive list of features and their corresponding values. These results are then integrated into the original clinical dataset, culminating in an enriched and context-aware augmented dataset. This enriched dataset not only contains new clinically relevant features but also ensures that each feature is accompanied by accurate, personalized values, enhancing the dataset's utility for subsequent analyses or predictive modeling.

Running Example (Phase III): DALL-M identifies that pleural effusion is commonly associated with the following new features: (1) Jugular Venous Distension (Yes/No), (2) Breath Sounds (Normal/Decreased/Absent), and (3) Chest Pain Severity (Mild/Moderate/Severe). The LLM assigns values based on the patient's known attributes. Given that the patient is 65 years old, has a temperature of 38.2 °C, and an oxygen saturation of 92%, DALL-M predicts: Jugular Venous Distension = Yes, Breath Sounds = Decreased, and Chest Pain Severity = Moderate. These features are added to the dataset, enriching the patient's clinical profile.

For more details, please refer to Appendix Figs. A.7 and A.6, where we outline how each component of the prompt interacts with the available data sources to yield high-quality, contextually appropriate results and present the general structure of our prompt formulation process and the methodology for assigning values.

While expert input refines the selection of clinically relevant queries, it is important to emphasize that DALL-M does not depend solely on expert-designed prompting for consistency. Instead, its robustness is achieved through a structured, three-phase pipeline, where (i) retrieval-augmented context processing ensures factual grounding, and (ii) structured augmentation maintains logical consistency. Even without expert-driven prompt engineering, the system retrieves and synthesizes clinically coherent information, demonstrating that performance improvements are inherent to its design rather than reliant on manual tuning.

4. Experimental setup

This section details the experimental setup used to evaluate the effectiveness of the DALL-M framework. We describe the datasets utilized, the clinical features selected for evaluation, and the specific hypotheses tested through a series of experiments.

4.1. Dataset

In this work, we utilize instances from the MIMIC-IV dataset [75], comprising radiographs, radiology reports from chest X-rays, and clinical data. We gathered 799 patient cases from these datasets, with each instance accompanied by corresponding ground-truth information on clinical features, facilitating the evaluation of the generated synthetic data. The labels used for evaluation in Section 5.2 are extracted from the REFLACX dataset [72], annotated by five radiologists. In the dataset, we have the following clinical features: (1) **temperature** (in Fahrenheit degrees), (2) **heartrate** (beats rate per minute), (3) **resprate** (breaths rate per minute), (4) **o2sat** (peripheral oxygen saturation as a percentage), (5) **sbp, dbp** (systolic and diastolic blood pressure, respectively, measured in millimeters of mercury (mmHg)). Additionally, the dataset incorporates ground truth labels indicating the presence/absence of lesions within the chest X-ray images of patients. We selected five distinct lesions for consideration: *atelectasis*, *consolidation*, *enlarged cardiac silhouette*, *pleural effusion*, and *pleural abnormality*. The rationale behind selecting these specific lesions stem from their

prevalence as the most frequently occurring conditions within the REFLACX dataset [72].

Although MIMIC-IV [75] contains over 40,000 patients, we curated a subset of 799 cases by ensuring that each sample contained all necessary modalities and labels for the study. The inclusion criteria required the presence of MIMIC-IV DICOM images, hospital module data (gender, age), emergency department triage features (such as vital signs), REFLACX lesion bounding box labels, and CheXpert labels [76]. This intersection resulted in a final dataset of 799 cases.

Additionally, while the five selected lesion classes are not perfectly balanced, we employed measures to mitigate potential biases, including data augmentation techniques and evaluation using Precision, Recall, and F1-score to ensure a fair assessment of model performance.

4.2. Experiments

This section outlines the experiments designed to evaluate the proposed DALL-M framework. These experiments aim to validate two key hypotheses central to the framework's ability to generate contextual values for *existing* features and extend its knowledge to create *new* clinically relevant features.

- **Hypothesis 1:** LLMs are capable of providing sufficient contextual information to generate accurate, patient-specific values for *existing* clinical features when provided with clinical patient reports.
- **Hypothesis 2:** LLMs can extend their knowledge to create entirely *new* clinically relevant features and their respective values, enhancing the original dataset.

To validate these hypotheses, we designed a series of experiments, each tailored to evaluate specific aspects of the LLMs' capabilities.

Validating Hypothesis 1: Can LLMs Generate Contextual Values for Existing Features? (see Section 5.1)

The first hypothesis explores whether LLMs can effectively generate values for existing clinical features based on available patient data. To test this, we conducted an experiment where we compared LLM-generated values with those produced by traditional methods, specifically Gaussian-based permutation techniques. No noise was added to image data. By comparing the outputs, we can assess whether LLMs are a reliable tool for generating patient-specific values that accurately reflect the clinical context.

Validating Hypothesis 2: Can LLMs Extend Their Knowledge to Generate New Features?

Once we validate that LLMs can generate accurate values for existing features, we move on to the second hypothesis, which explores whether LLMs can extend their knowledge to create new clinically relevant features and generate their values. This requires a more sophisticated experiment design, which we carried out in the following stages:

- **Stage 1: Analyzing the Quality of Information Retrieval** (see Section 5.2). The first step in validating Hypothesis 2 involves understanding what data sources contribute to generating high-quality, contextually relevant information. To do this, we used data from two complementary sources: Wikipedia and Radiopaedia. We tested which of these sources (Wikipedia, Radiopaedia or both) give the LLM access to a more comprehensive clinical knowledge base, allowing it to generate better answers for both general and specialized queries (see Phase I in Fig. 1). Other sources, such as medical textbooks and clinical databases (e.g., UpToDate, PubMed, and professional guidelines), contain high-quality expert-reviewed content. However, their integration into our framework is limited by copyright and access restrictions, preventing their direct use in automated retrieval pipelines.

- **Stage 2: Generating Prompts for New Feature Generation (see Section 5.3).** After understanding which data sources lead to better information retrieval, the next step was to design prompts that guide the LLM in generating new clinical features. Based on a set of questions given to radiologists, these prompts are tailored to retrieve information from Radiopedia and Wikipedia and to *extract novel features from the LLM answers*. This extends the LLM's understanding beyond existing features, allowing it to synthesize entirely new, contextually relevant features from the knowledge it has gained in Stage 1 (see Phase II in Fig. 1). In this stage, we designed retrieval-specific prompts that guide the LLM in generating context-aware features. Unlike approaches that fine-tune prompts through systematic training, we employed a trial-and-error method to iteratively refine the prompts based on observed model behavior.
- **Stage 3: Populating the Generated Features with new Feature Values (see Section 5.3).** After identifying new features, we used the RAG framework to generate accurate values for each feature using the contextual knowledge established in previous phases. The RAG framework first retrieves relevant information from the Augmented Clinical Knowledge corpus, and then the generator model synthesizes this information to produce clinically coherent values. These generated values are then integrated into the original dataset, creating a richer, contextually enhanced dataset that aligns with the patient's clinical profile and is ready for downstream analysis.
- **Stage 4: Evaluating the Effectiveness of the Newly Generated Features (see Section 5.3).** We carried out a comprehensive set of experiments to evaluate the new features generated by the LLMs. First, we **augmented the dataset with the newly generated features** and tested their impact on several machine learning models, including Decision Trees, Random Forest, XGBoost, and TabNet. The models were evaluated using metrics such as accuracy and AUC to determine the added value provided by the new features. We also performed **an ablation study** to further investigate the contribution of the Augmented Clinical Knowledge from Phase II. By isolating the ACK's influence on the LLM's ability to generate valuable features, we assessed its critical role in the feature-generation process. Finally, we conducted **a feature importance analysis** to identify which of the newly generated features had the greatest relevance to clinical outcomes. These findings were subsequently reviewed by medical experts to ensure the clinical validity and relevance of the generated features. In evaluating the effectiveness of the newly generated features, we initially report accuracy and AUC as general performance indicators. However, we recognize that in clinical applications, false negatives are particularly concerning, as missing relevant clinical information could impact diagnosis and treatment decisions. To address this, our results emphasize precision and recall, as they provide a more meaningful assessment of feature validity. Precision ensures that newly generated features are clinically relevant, while recall minimizes the risk of missing critical information. This dual focus allows us to balance overall model performance with medical reliability.

4.3. Parameters and model setup

In our experiments, no fine-tuning was performed, as we leveraged a Retrieval-Augmented Generation approach that directly incorporates medical sources, eliminating the need for model-specific adaptation. Since each LLM has its own internal mechanisms for interpreting queries, there is no standardized method for prompt formulation. To ensure a fair evaluation, we used a consistent query formulation across all models without modifying LLM-specific prompt structures, preventing biases in prompt optimization and ensuring that each model's

response reflects its native ability to retrieve and synthesize clinical knowledge. Additionally, we set the LLM temperature to 0.1 to prioritize deterministic and clinically consistent outputs, as higher temperatures (e.g., 0.7 or above) occasionally led to clinically implausible outputs, such as symptoms unrelated to the patient's condition. Given the sensitive nature of medical data augmentation, this choice was guided by preliminary tests, clinical requirements, and consistency considerations, ensuring that generated outputs align closely with retrieved domain-specific knowledge from Radiopaedia and the patient context.

In the downstream evaluation of our augmented dataset, we employed default parameters for each machine learning model to ensure consistency and reproducibility. For Decision Trees and Random Forests, we used default settings, with Decision Trees having a maximum depth of 10 and Random Forests consisting of 100 trees, both utilizing the Gini impurity function for node splitting. XGBoost was configured with a learning rate of 0.01, a max depth of 6, and 100 estimators, using the binary:logistic setting to accommodate the binary nature of our classification tasks. For the TabNet model, we leveraged the PyTorch implementation from Dreamquark AI with its preset parameters, including an optimizer learning rate of 0.02, a batch size of 1024, and a decision dimension of 8 with 2 steps in the decision process.

All experiments used high-performance computational resources, including NVIDIA 4090 24 GB GPUs and Intel i9-13900K CPUs. These setups allowed for efficient training and testing of the machine learning models and rapid processing of the large datasets required for the experiments.

4.4. Downstream task parameters

In the downstream evaluation of our augmented dataset, we employed default parameters for each machine learning model to ensure consistency and reproducibility. For Decision Trees and Random Forests, default settings were applied, Decision Trees having a maximum depth of 10 and Random Forests consisting of 100 trees, both using the Gini impurity function for node splitting. For XGBoost, default parameters included a learning rate of 0.01, a max depth of 6, and 100 estimators, with the model set to 'binary:logistic' to accommodate the binary nature of our classification tasks. For the TabNet model, we used the PyTorch implementation from Dreamquark AI with its preset parameters. These include an optimizer learning rate of 0.02, a batch size of 1024, and a decision dimension of 8 with 2 steps in the decision process.

5. Results

5.1. Experiment I: Can LLMs generate contextual values for existing features

This experiment evaluated the potential of LLMs to accurately generate synthetic values for pre-existing clinical features within the MIMIC-IV dataset. The criticality of this task cannot be overstated, given the propensity of LLMs to produce "hallucinations" or inaccuracies that, within a medical context, could lead to significant ramifications for patient care. To mitigate these risks and ensure the generation of clinically relevant values, we employed a structured approach involving generating a prompt (Fig. 2). We set the temperature parameter to 0.1 and applied eight clinical features from the MIMIC-IV dataset (oxygen saturation, temperature, gender, age, heart rate, o2sat, systolic (sbp), and diastolic blood pressure (dbp)) to 799 patient cases.

We analyzed the capabilities of different types of LLMs to generate clinically relevant synthetic data. We tested nine different LLMs ranging from medical domain-specific models (such as BioGPT [68], ClinicalBERT [77], BioClinicalBERT [78], and Medtiron [79]) to more general models (Mistral [80], Zephyr [81], Llama2 [82], GPT-3.5, and GPT-4).

```

INDICATION: Central venous line placement.
TECHNIQUE: Frontal chest radiograph.
COMPARISON: Chest radiograph 12:42 today.
FINDINGS: A right subclavian catheter has been placed in the interim. The catheter terminates at the
confluence of the brachiocephalic vein and superior vena cava and, if indicated, could be
advanced 3.7... cm for termination within the low SVC. There is no pleural effusion or
pneumothorax. The cardiac silhouette remains mildly enlarged. There is no focal airspace
consolidation worrisome for pneumonia. High density material is again seen in the paritally
imaged colon in the left abdomen. Cholecystectomy clips are noted. There are carotid
calcifications, left greater than right.
AGE: 69.
GENDER: Female.
LESIONS: Enlarged cardiac silhouette.
Oxygen Saturation (%): [MASK].

```

Fig. 2. Example of a prompt used in Experiment I for generating feature values for Oxygen Saturation (%).

Table 2
Mean Square Error for normalized clinical features.

Model	Temperature	Heartrate	Resprate	o2sat	sbp	dbp	Mean
BioGPT [68]	6.124	1.733	108.269	7.799	3.628	1.836	21.565
ClinicalBERT [77]	5.623	2.663	22.880	2.298	6.086	1.443	6.832
BioClinicalBERT [78]	4.204	1.578	7.230	1.466	5.007	1.661	3.524
Mistral [80]	3.635	1.920	5.842	7.824	2.284	3.056	4.093
Zephyr [81]	2.144	1.541	3.060	4.146	1.770	1.700	2.393
Llama2 [82]	1.318	1.559	2.617	2.387	1.494	3.741	2.186
Meditron [79]	2.420	2.057	4.858	6.359	2.205	4.071	3.662
GPT-3.5	1.325	1.308	1.159	3.403	1.264	2.338	1.799
GPT-4	1.966	1.344	1.298	1.722	1.251	1.705	1.548
MMTF (supervised)	0.003	0.008	0.040	0.009	0.006	0.013	0.013
Gaussian (baseline)	2.074	2.085	2.027	1.987	2.092	1.934	2.033

Fig. 2 presents an example of a prompt to generate values for existing features based on the patient's clinical context.

To quantify the deviation of LLM-generated values from actual clinical data, we employed the Mean Squared Error (MSE) metric, comparing the generated values against the ground truth for each feature. We use this metric to indicate the LLM's accuracy and reliability in synthesizing clinical data.

We used a simple Gaussian distribution to generate clinical agnostic feature values as a baseline. To compare the performance of LLM-generated clinical feature values with a multimodal approach, we implemented a Multimodal Transformer (MMTF) inspired by ViT [83]. The MMTF was trained using chest X-ray images and corresponding structured clinical features from the MIMIC-IV dataset. The model was optimized using a cross-entropy loss function and fine-tuned over 200 epochs with an Adam optimizer (learning rate = 0.0001). Evaluation was conducted using the same Mean Squared Error metric applied to LLM-generated values, ensuring a direct comparison between methods.

The experiment was structured to run each query through the multimodal vision transformer, mirroring the approach taken with the LLMs to ensure comparability. Table 2 reports the results using MIMIC-IV, where GPT-4 provides the best performance using mean MSE (i.e., average overall features).

5.2. Experiment II: Analyzing the quality of information retrieval

In Experiment II, we evaluated the effectiveness of information retrieval and data storage mechanisms by comparing LLMs' performance using data from Radiopaedia and Wikipedia. The experiment focused on two main aspects: the information sources and the data representation types.

Using the best LLM from Experiment I, GPT-4, we conducted queries using three configurations: Wikipedia alone, Radiopaedia alone, and combined. Due to the high costs of using GPT-4, we limited our experiments to a single query: "Atelectasis". The following conditions were tested: (a) all documents from the first result page of Wikipedia

(approximately 10 documents); (b) all documents from the first result page of Radiopaedia (approximately 20 documents); (c) the top result from Wikipedia (1 document); (d) the top result from Radiopaedia (1 document); and (e) all documents from the first result pages of both Wikipedia and Radiopaedia (approximately 30 documents total). For each setting, we tested three types of data representation: NetworkX, Neo4j Vector databases, and Neo4j Graph databases. The goal was to retrieve information for the domain questions outlined in Phase II: (1) What is Atelectasis?, (2) What are the symptoms of Atelectasis?, (3) What can cause Atelectasis?, (4) What clinical features do radiologists need for diagnosing Atelectasis?, and (5) What are the radiographic features of Atelectasis?

The most effective retrieval configuration was determined by comparing different data sources (Wikipedia, Radiopaedia, or both) and storage methods (NetworkX, Neo4j Vector, and Neo4j Graph databases). The assessment was based on the completeness, accuracy, and contextual relevance of the retrieved information in answering domain-specific questions (e.g., symptoms, causes, and radiographic features of atelectasis). The best-performing setup, using both Wikipedia and Radiopaedia stored in a Neo4j Vector database, provided the most comprehensive and clinically meaningful responses.

Our experiments show that the best data representation is the Neo4jVector (see Table 3 for an example of LLM outputs). This is due to its ability to capture and retrieve semantic relationships within high-dimensional data efficiently. The Neo4j Vector database outperformed other methods by providing more contextually relevant and accurate responses to the domain questions. Its advanced indexing and search capabilities allowed for better integration of diverse information sources, ultimately enhancing the LLM's ability to generate high-quality, clinically relevant features. This demonstrates that leveraging a combination of broad and specialized knowledge stored in a semantically rich vector format significantly improves the performance of the LLM in clinical data augmentation tasks.

Table 3

GPT-4 Responses based on Wikipedia and Radiopaedia: Synthesis of information retrieved from the top search results for query “Atelectasis”. The highlighted terms indicate areas where additional clinical context was retrieved to guide the feature generation process. While these terms alone do not directly translate into new features, they help DALL-M retrieve relevant medical knowledge, which is then used to synthesize new, context-aware features.

Question	LLM Answer
Networkx	
1	Atelectasis is a condition in which small areas of the lung collapse due to adjacent compression, passive atelectasis, or dependent atelectasis. It is usually seen on chest X-rays as small volume linear shadows and CT is often helpful to determine the cause or confirm that there is no proximal obstruction.
2	Atelectasis is usually asymptomatic , but some people may experience shortness of breath, chest pain, and coughing
3	Atelectasis can be caused by adjacent compression, passive atelectasis, and dependent atelectasis
4	Radiologists need to look for small areas of collapsed lung, small volume linear shadows, and adjacent compression on chest X-rays to diagnose Atelectasis. CT scans may also be helpful to determine the cause or confirm that there is no proximal obstruction
5	Atelectasis is usually seen on chest X-rays as small volume linear shadows
Neo4jVector	
1	Atelectasis is the collapse or closure of a lung resulting in reduced or absent gas exchange . It is usually unilateral, affecting part or all of one lung. It is a condition where the alveoli are deflated down to little or no volume. It can be caused by normal exhalation or various medical conditions. Atelectasis is not synonymous with a pneumothorax , although a pneumothorax can cause atelectasis. It can occur as a post-operative complication or as a result of surfactant deficiency . It is a common finding in chest X-rays and other radiological studies. The most common cause is post-surgical atelectasis, which develops in 75%–90% of people undergoing general anesthesia for a surgical procedure. Other causes include pulmonary tuberculosis, obstruction of the airway, poor surfactant spreading, and compression of the lung by a space-occupying lesion. Clinically significant atelectasis is visible on chest X-ray and may require further imaging or bronchoscopy for diagnosis. Treatment involves managing the underlying cause
2	The symptoms of Atelectasis may include cough (not prominent), chest pain (not common), breathing difficulty (fast and shallow), low oxygen saturation, pleural effusion (transudate type), cyanosis (late sign), and increased heart rate . However, it is important to note that atelectasis can also be asymptomatic
3	Atelectasis can be caused by various factors, including post-surgical complications, surfactant deficiency, pulmonary tuberculosis, blockage of the airway, compression from external factors, poor surfactant spreading, and certain risk factors such as type of surgery, obesity, high oxygen, and presence of chronic obstructive pulmonary disease or asthma . It can also occur due to passive atelectasis or adhesive atelectasis
4	Radiologists need to look for small volume linear shadows on chest X-rays, usually peripherally or at the lung bases, to diagnose atelectasis. They can also use CT scans to accurately depict the involved lung and determine the cause. Additionally, the presence of the crow feet sign , which appears as linear bands radiating from a mass into adjacent lung tissue, can be a characteristic feature of round atelectasis. The clinical presentation of atelectasis can vary depending on the extent and speed of development, but common symptoms include breathlessness, cough, and diminished breath sounds . The underlying causes of atelectasis can include adjacent compression, passive atelectasis, and dependent atelectasis . Treatment involves managing the underlying cause and promoting deeper breathing if hypoventilation is the cause.
5	Radiographic features of atelectasis include small volume linear shadows on chest X-rays, usually peripherally or at the lung bases. The underlying cause may also be visible. CT chest allows for a more accurate depiction of the involved lung and can help determine the cause. Round atelectasis can be seen on CT as linear bands radiating from a mass into adjacent lung tissue, resembling the feet of a crow . Atelectasis can also be classified based on underlying mechanism (resorptive, passive, compressive, cicatrization, adhesive, gravity dependent) or morphology (linear, round, lung, lobar, segmental, subsegmental). The radiographic features vary depending on the type and mechanism of atelectasis.
Neo4jGraph	
1	Atelectasis is a condition where part or all of one lung collapses or closes, resulting in reduced or absent gas exchange . It can be caused by normal exhalation or various medical conditions. The alveoli, which are small air sacs in the lungs, become deflated down to little or no volume. Atelectasis is also known as a collapsed lung.
2	I'm sorry, but I don't have the information to answer your question. It would be best to consult a medical professional for accurate information on the symptoms of Atelectasis
3	There are several factors that can cause atelectasis. Some common causes include blockage of the airways due to mucus, a foreign object, or a tumor , as well as lung diseases such as chronic obstructive pulmonary disease (COPD) or pneumonia . Additionally, certain medical procedures, such as anesthesia or prolonged bed rest , can also contribute to the development of atelectasis. It is important to consult with a healthcare professional for a proper diagnosis and treatment plan.
4	To diagnose Atelectasis, radiologists typically look for certain clinical features. These may include signs such as lung collapse, airway obstruction, and reduced lung volume . Additionally, they may also consider the presence of symptoms such as shortness of breath, coughing, and chest pain . By evaluating these clinical features, radiologists can make an accurate diagnosis of Atelectasis
5	I'm sorry, but I don't have the information to answer your question about the radiographic features of Atelectasis.

5.3. Experiment III: Evaluating the effectiveness of the newly generated features

Building on the initial assessment of LLMs to generate realistic values for existing clinical features, this experiment evaluates DALL-M's capability to create new, clinically relevant features to enhance the dataset comprehensively. This investigation involved processing 799 patient case data points, originally comprising only eight features, to include new features and their corresponding synthetic values, as established in Experiment I. We termed this method *Augmented*, containing 78 features. Furthermore, we consulted medical experts to leverage clinical expertise and enrich the dataset. They reviewed the newly generated features and recommended including specific clinical features based on their professional experience. The subsequent generation of values, denoted as *Augmented with Expert Input*, yielded 91 features.

To rigorously assess the contribution of these augmented features to the dataset, three distinct analytical approaches were employed: (1) We

evaluated the augmented dataset using various ML models, including Decision Trees, Random Forest, XGBoost, and TabNet. This analysis used metrics such as accuracy and area under the curve (AUC); (2) An ablation study was conducted to discern the impact of incorporating ACK into the prompt generation phase on the overall effectiveness of the augmentation process; (3) The ML models employed underwent a feature importance analysis. Our medical experts then reviewed the results to validate the clinical relevance of the identified features.

We documented the outcomes in Table 4. In the table, The “#Rel. Features” column in Table 3 represents the number of features that were found to be important for each machine learning method. We measured feature importance across the generated and original features, and this column reflects how many features significantly contributed to the predictive performance of each algorithm. An example of feature importance distribution for Decision Trees, Random Forests, and XGBoost is presented in Fig. 3 and TabNet in Fig. 4, with expert commentary on the identified features' clinical validity discussed in

Table 4

Performance of different feature sets on REFLACX dataset [72] using CheXpert [76] labels as groundtruth. The bold values represent the best-performing models, while the underlined values correspond to the second-best. (#Rel. F. = Relevant Features)

Feature set	Model	Accuracy	AUC	Precision	F-1	Recall	P-value	#Rel. F.
Original # feat: 8	DecisionTree	0.6676	0.5666	0.367	0.3591	0.352	0.3021	0
	RandomForest	0.7484	0.554	0.6091	0.2285	0.1408	0	0
	XGBoost	0.7108	0.5389	0.3953	0.2411	0.1735	0	8
	TabNet	0.6568	0.5250	0.3117	0.2743	0.2449	0.5000	6
Augmented # feat: 78	DecisionTree	0.7319	0.6124	0.4935	0.4139	0.3582	0.0222	6
	RandomForest	0.7749	0.6089	<u>0.7067</u>	0.3752	0.2561	0	3
	XGBoost	0.7865	0.6687	0.6508	0.5093	0.4184	0.0001	49
	TabNet	0.7027	0.5856	0.4231	0.375	0.3367	0.0293	22
Augmented with Expert input # feat: 91	DecisionTree	0.7246	0.6244	0.4767	0.4413	0.4112	0.2322	5
	RandomForest	0.7811	0.6177	0.7362	0.3953	0.2704	0	4
	XGBoost	<u>0.7838</u>	<u>0.6636</u>	0.6452	<u>0.5000</u>	0.4082	0.0002	<u>56</u>
	TabNet	0.7243	0.6232	0.4762	0.4396	0.4082	0.3974	43
Ablation No Prior Knowledge # feat: 91	DecisionTree	0.7127	0.6182	0.4549	0.4349	<u>0.4173</u>	0.6686	1
	RandomForest	0.7705	0.6174	0.6483	0.4022	0.2918	0	1
	XGBoost	0.7486	0.6365	0.5342	0.4561	0.398	0.4082	47
	TabNet	0.6541	0.5754	0.3636	0.3846	0.4082	0.6799	59

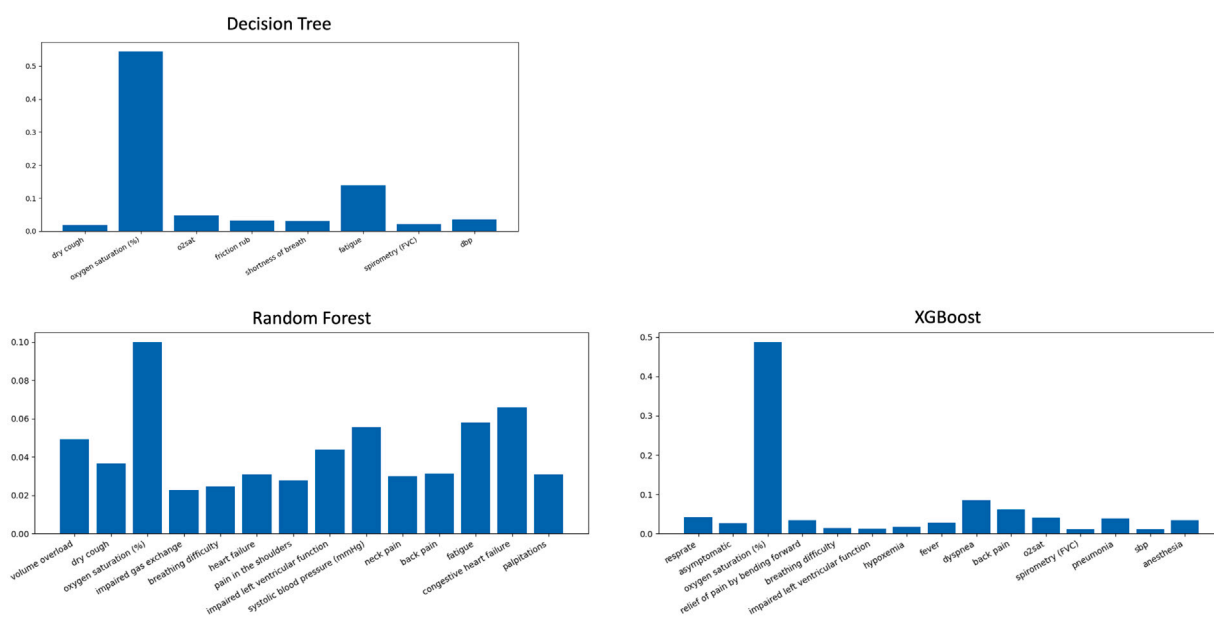


Fig. 3. Example of feature importance distribution for enlarged cardiac silhouette using Decision Trees, Random Forests and XGBoost.

DALL-M synthesized 70 new clinical features, later expanded to 91 after expert input. These features include lesion-specific attributes (e.g., pleural effusion severity, nodule shape descriptors), comorbidity indicators (e.g., history of chronic obstructive pulmonary disease), and additional symptom-related variables inferred from radiology reports. In the Appendix, Table B.5 provides a detailed breakdown of these synthesized features.

Fig. 5 presents the Pearson correlation coefficients between the feature importance scores of four machine learning models (XGBoost, Random Forest, Decision Tree, and TabNet) and the clinical evaluations provided by radiologists. Each panel corresponds to a specific condition – Atelectasis, Consolidation, and Enlarged Cardiac Silhouette – with separate plots for features rated as high relevance (left column) and low relevance (right column). The bars represent the average correlation with two clinical assessment criteria: clinical plausibility and clinical relevance for the set of features that were ranked least and highly important by two radiologists.

6. Discussion

Our experiments with the DALL-M framework highlight the accuracy of LLMs in generating synthetic clinical values and creating new

clinically relevant features. The precision of models like GPT-4 and the enhancements from integrating expert insights are key focuses. We examine the impact of these advancements on predictive modeling, clinical decision-making, and patient care and contemplate the future role of AI in healthcare. The main contributions are as follows:

(1) LLMs Can Generate Clinically Relevant Synthetic Values: Experiment I demonstrates the potential of LLMs, particularly GPT-3.5 and GPT-4, to accurately generate synthetic values for existing clinical features. Using structured prompts, our evaluation through MSE suggests the models’ ability to produce values close to real clinical data, with GPT-4 showing the best performance among LLMs.

(2) Multimodal Approaches Offer Further Improvements: Integrating LLMs with other data modalities, such as images, could enhance the accuracy and relevance of generated clinical feature values. Although we explored this only in Experiment I, further research could incorporate these models in Phase III.

(3) General LLMs perform better than Domain-Specific ones: Experiment I shows that general LLMs outperform domain-specific LLMs in generating synthetic clinical values. This may be due to several factors: (i) general LLMs are trained on a vast and diverse dataset, including ample medical literature, which may surpass the performance

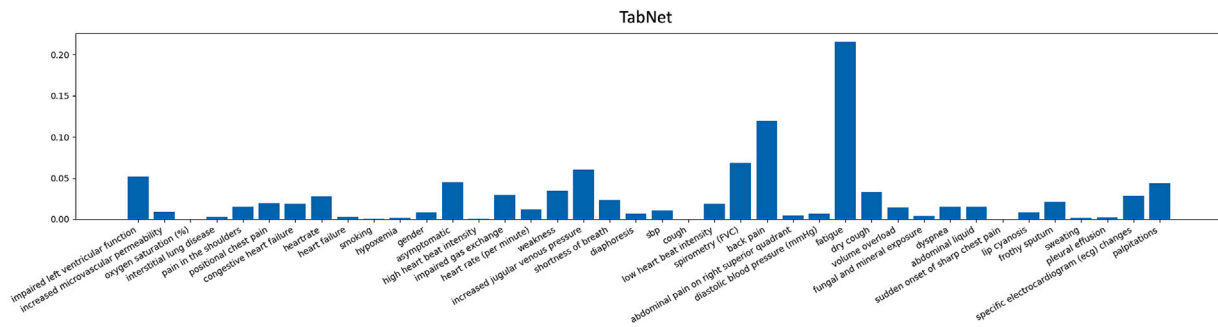


Fig. 4. Example of feature importance distribution for enlarged cardiac silhouette using TabNet.

of models trained on narrower datasets, and (ii) general LLMs are updated regularly, leading to up-to-date information.

(4) DALL-M Improves Model Performance Despite the Complexity of Clinical Data: Experiment III demonstrates DALL-M's ability to augment existing datasets with synthetic values and expand them with entirely new, clinically relevant features. Without medical expert input, DALL-M extended the eight original features to 78 features, significantly enhancing model performance. Despite the inherent challenges of working with sparse and highly structured clinical tabular data, our results show a **25 percentage point increase in Precision and Recall** for XGBoost and a **16.5 percentage point improvement in F1-score**. These gains highlight the potential of DALL-M in addressing baseline performance limitations, particularly in real-world medical datasets that are often incomplete or imbalanced. While Accuracy is a widely used metric, it can be misleading in medical datasets with class imbalances. In contrast, Precision, Recall, and F1-score are more clinically relevant, as they directly address the trade-off between false positives and false negatives. High Recall is crucial to avoid missing critical cases, while high Precision reduces unnecessary interventions, making these metrics more suitable for evaluating medical diagnostic tasks. Additionally, data augmentation with expert input generated 13 additional features, totaling 91 features, leading to a further 3% increase in Precision for Random Forests. These findings emphasize the practical benefits of context-aware feature augmentation in improving predictive modeling in healthcare. XGBoost consistently outperformed other models across most metrics, reaffirming the well-established effectiveness of tree-based models in handling tabular data [84,85].

(5) Ablation studies confirm - Augmented Clinical Knowledge is vital: The ablation study further confirms that DALL-M's effectiveness is a result of its structured augmentation pipeline rather than reliance on expert-crafted prompts. While expert input refines the prompting process, its role is to enhance retrieval efficiency rather than inject external knowledge. The system's ability to generate high-quality synthetic values and new features primarily derives from its retrieval-augmented framework, which anchors responses in trusted medical sources. This ensures that clinical consistency is achieved by design, independent of manual adjustments to prompts. As shown in Table 4, the absence of augmented clinical knowledge results in a notable performance decline across all models. This underscores the role of augmented clinical knowledge in enhancing the ability of LLMs to produce reliable and precise synthetic values. Nevertheless, our approach requires ethical scrutiny for future iterations, to ensure accuracy, transparency, and mitigate biases [86–88].

(6) DALL-M Excels when Compared with Traditional Clinical Data Augmentation Methods: Traditional augmentation techniques – such as SMOTE, GANs, and Gaussian noise injection – primarily manipulate existing data without introducing new, clinically meaningful information. SMOTE interpolates between existing samples, failing to enrich datasets with novel insights. GAN-based approaches generate synthetic data but often lack interpretability and can introduce artifacts. Gaussian perturbation methods inject randomness, risking

implausible clinical relationships. DALL-M surpasses these methods by generating entirely new, context-aware features grounded in verified medical knowledge through RAG. This ensures medical consistency, interpretability, and enhanced predictive power. By integrating expert-driven contextualization with advanced AI techniques, DALL-M fundamentally redefines how clinical datasets are enriched, moving beyond simple augmentation to intelligent knowledge-driven feature expansion.

(7) Alignment of Machine Learning Models with Radiologists' Clinical Assessments: The correlation analysis indicates that Random Forest aligns most closely with radiologists' assessments of clinical relevance, particularly for high-relevance features in Consolidation and Enlarged Cardiac Silhouette. This suggests that Random Forest effectively identifies diagnostically significant features, likely due to its ensemble approach, which enhances robustness and reduces overfitting. In contrast, TabNet tends to overestimate feature importance, especially for Enlarged Cardiac Silhouette, resulting in weaker correlations with radiologists' evaluations. This overestimation suggests that TabNet may capture broader patterns but struggles to differentiate between clinically critical features and less relevant ones, highlighting the importance of aligning machine learning feature selection with expert clinical judgment for improved interpretability and diagnostic support.

7. Challenges, limitations and future work

The DALL-M framework shows promise for enhancing clinical data, but faces several key challenges that need addressing.

While our results demonstrate the effectiveness of DALL-M, several challenges must be acknowledged for real-world deployment. First, LLM-generated data must maintain clinical consistency, as synthetic values may not always align with real-world patient distributions. Second, bias propagation remains a concern—LLMs trained on imbalanced medical datasets may reinforce existing disparities, leading to skewed predictions. Third, unlike real patient records, synthetic data lacks direct ground truth validation, making rigorous quality control essential. Additionally, the use of LLMs in clinical augmentation raises regulatory and ethical challenges, particularly regarding compliance with GDPR, and medical AI guidelines. Finally, LLMs may introduce hallucinations, where fabricated correlations or features could mislead predictive models if not carefully filtered. Future research should explore hybrid methods that combine LLM augmentation with domain expert validation to mitigate these risks.

The choice between specialized and general-purpose language models presents another dilemma. While GPT-4 currently outperforms medical-specific models, its effectiveness in specialized medical sub-fields remains uncertain. Additionally, the high computational costs of running advanced models like GPT-4 create practical barriers to widespread clinical adoption. Ethical considerations also loom large. Though we use synthetic data to protect patient privacy, we must carefully monitor for potential biases and prevent misuse of generated clinical information. Moving forward, addressing these challenges

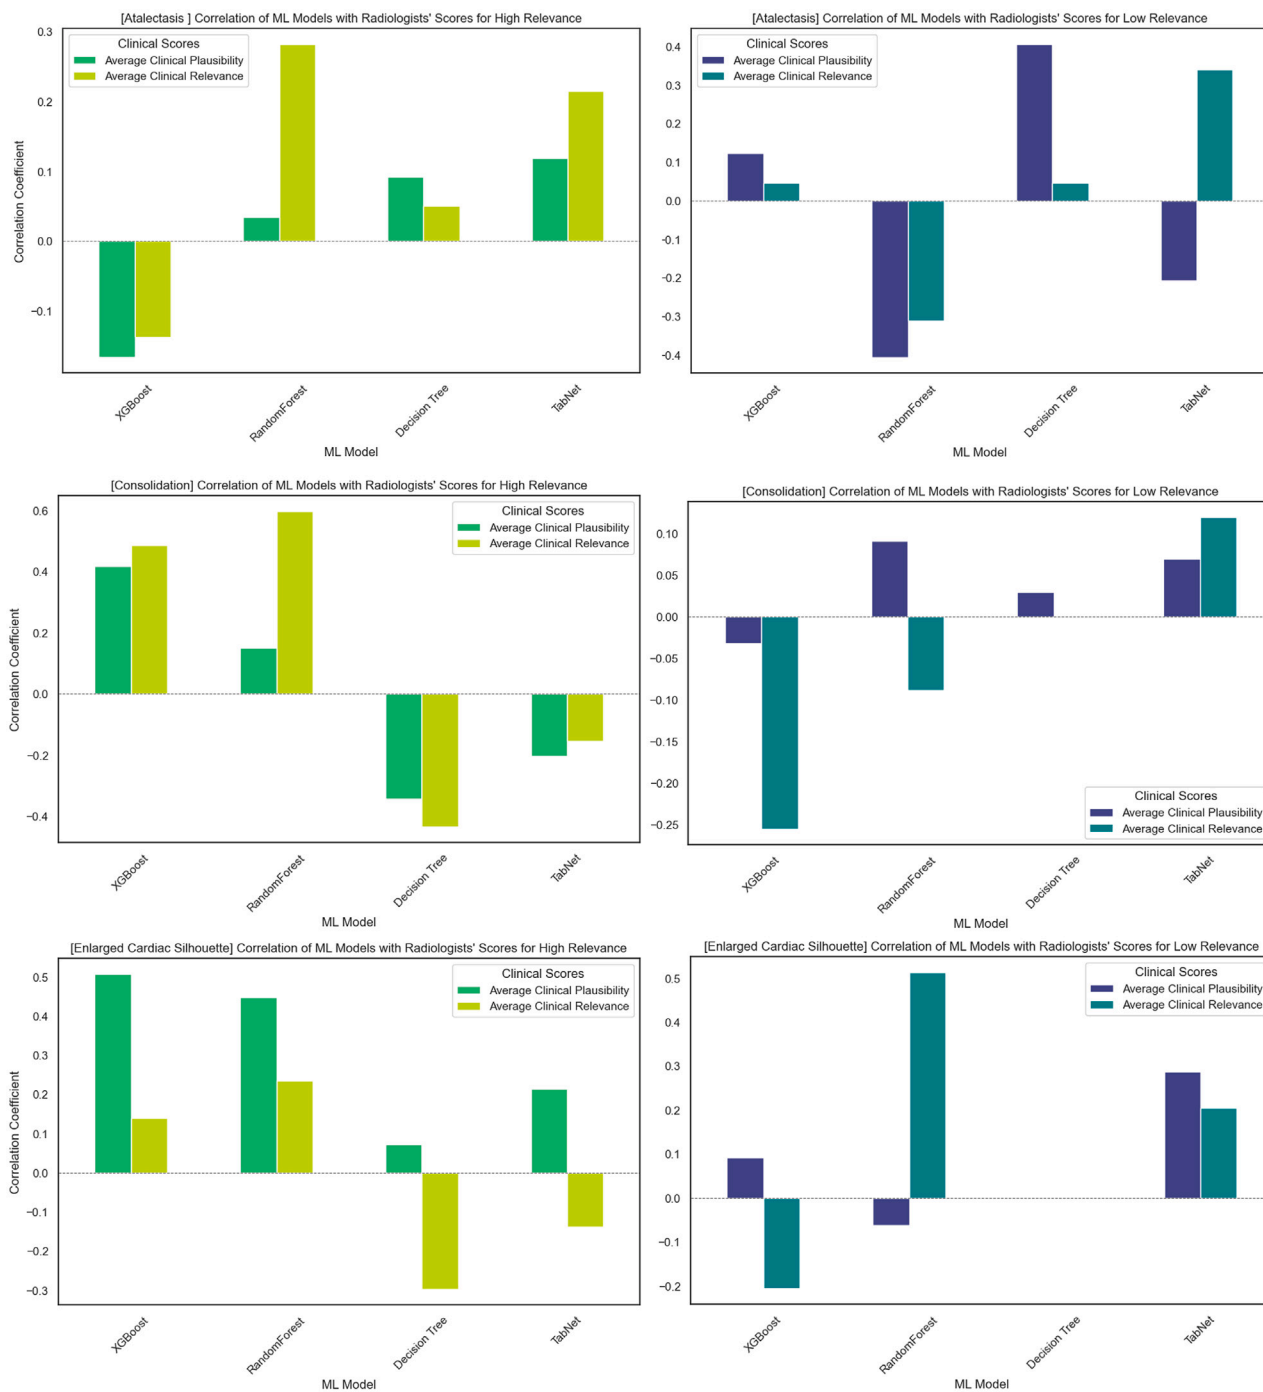


Fig. 5. Correlation of machine learning models with Radiologists' clinical scores for high and low relevance features.

while maintaining strict ethical guidelines will be crucial for successful clinical implementation.

For future work, we aim to explore bias mitigation strategies to ensure fairness and reliability in synthetic data generation. Techniques such as adversarial debiasing, reweighting methods, and fairness-aware augmentation will be investigated to address potential disparities introduced by LLM-generated clinical data. Additionally, we plan to expand DALL-M's capabilities by integrating multimodal clinical data, including radiology images and electronic health records (EHRs). By leveraging multimodal fusion techniques, we can enhance the contextual depth of synthetic feature generation, improving its alignment with real-world medical scenarios.

8. Conclusion

This study proposes two distinct strategies for augmenting clinical datasets. The first approach addresses the issue of missing values by generating realistic synthetic values using trained LLMs and multimodal transformers for supervised learning. This method ensures the dataset is complete and robust, allowing machine learning models to learn more effectively.

The second approach focuses on generating clinically relevant features by leveraging the combined power of LLMs and medical experts' domain knowledge. This method goes beyond simply filling in gaps; it creates new, valuable data points that enhance the overall quality of

the dataset. By incorporating expert insights, we ensure the generated features are statistically sound and clinically meaningful.

Our results demonstrate that using the augmented dataset significantly improves model performance across various machine learning models. The features identified through our augmentation strategies are validated as clinically relevant, showing potential to enhance diagnostic accuracy and patient care. To the best of our knowledge, DALL-M is the only method capable of generating new context-aware features in tabular data. This capability is particularly impactful in fields where data is scarce, as it can significantly enhance the quality and quantity of available data, leading to better model performance and more informed decision-making.

CRediT authorship contribution statement

Chihcheng Hsieh: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Catarina Moreira:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Isabel Blanco Nobre:** Validation, Methodology, Formal analysis. **Sandra Costa Sousa:** Validation, Methodology, Formal analysis. **Chun Ouyang:** Writing – review & editing, Writing – original draft, Supervision. **Margot Brereton:** Writing – review & editing, Writing – original draft, Supervision, Resources. **Joaquim Jorge:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Investigation, Funding acquisition, Conceptualization. **Jacinto C. Nascimento:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The work reported in this article was partially supported by the Center of Data Science at Queensland University of Technology and was partially supported under the auspices of the UNESCO Chair on AI & VR by national funds through Fundação para a Ciência e a Tecnologia with references DOI:10.54499/UIDB/50021/2020, DOI:10.54499/DL57/2016/CP1368/CT0002 and 2022.09212.PTDC (XAVIER project).

Appendix A. Prompt for synthetic clinical feature value generation

Fig. A.7 illustrates a concrete example of the final prompt used in the Synthetic Clinical Feature Value Generation Module. This prompt leverages the patient's clinical report, extracted contextual information, and the expert knowledge synthesized in the Augmented Clinical KnowledgeModule to generate clinically relevant synthetic values for various features. The LLM is instructed to take on the role of an experienced radiologist, using its domain-specific knowledge to make predictions for specific clinical features such as “breathing difficulty” or “pleural effusion”.

The prompt begins by providing the patient's detailed report, including radiological findings and clinical impressions. Next, it outlines the relevant prior knowledge regarding the condition (in this case, atelectasis), covering its symptoms, potential causes, and diagnostic criteria. The LLM is then asked to assess whether the patient's clinical presentation aligns with known patterns of atelectasis and other relevant conditions.

The final section of the prompt structures the LLM's responses by using a template with binary feature values (e.g., True/False) accompanied by a rationale. This ensures that the generated values are not only accurate but also explainable, with each prediction supported by clinical reasoning. For example, the LLM might indicate “breathing difficulty: True, because the patient's report mentions fast and shallow breathing, which is a typical symptom of atelectasis”.

This structured and context-aware prompt enables the LLM to generate patient-specific values for existing features and introduce new clinically relevant features, enhancing the overall quality of the dataset for downstream analyses and machine learning applications.

Appendix B. Augmented feature list

A comprehensive list of the 70 new features generated by DALL-M is provided, including their clinical category, average relevance score, and novelty assessment as determined by radiologists.

Data availability

The DALL-M source code is publicly available at: <https://github.com/ChihchengHsieh/DALL-M>.

```
System Prompt: You are an experienced radiologist with more than 30 years of experience. You are the
most respected radiologist in the world. You are examining a patient with the following report:

<Patient's Radiology Report>
=====

According to your prior knowledge:

<Augmented Clinical Knowledge Output>
=====

User Prompt: Make a prediction based on your prior knowledge and on the patient's report. Please
check if your prediction is in accordance with your prior knowledge. Ensure your answers are
using the following template.
<new feature 1>: {True / False}, because...
<new feature 2>: {True / False}, because...
...
<new feature n>: {True / False}, because...
```

Fig. A.6. General template of the prompt used in the Synthetic Clinical Feature Value Generation Module with n new features generated by the Augmented Clinical Knowledge Module. For a concrete example, please refer to Appendix Fig. A.7.

You are an experienced radiologist with more than 30 years of experience. You are the most respected radiologist in the world. You are examining a patient with the following report:

EXAMINATION: PA AND LATERAL CHEST RADIOGRAPHS INDICATION: yearold male with weakness. Evaluate for pneumonia. TECHNIQUE: PA and lateral chest radiographs COMPARISON: Multiple prior chest radiographs FINDINGS: Compared with the previous examination there is increased diffuse interstitial opacities with a linear consolidation in the right lower lung compatible with atelectasis. There is also a focal opacity in the left for lower lung and retrocardiac region with associated small pleural effusion better seen in the lateral view. Calcified granulomas are redemonstrated more prominently in the left apex. Moderate cardiomegaly is stable. Rightward deviation of the trachea is also unchanged. There is no pneumothorax. IMPRESSION: Patchy opacities in the left lower lung and retrocardiac region may represent left lower lobe pneumonia on the background of interstitial pulmonary edema. LESIONS: Enlarged cardiac silhouette Atelectasis. AGE: 93. GENDER: Male.

=====
According to your prior knowledge:

The symptoms associated with atelectasis may include cough (not prominent), chest pain (not common), breathing difficulty (fast and shallow), low oxygen saturation, pleural effusion (transudate type), cyanosis (late sign), and increased heart rate. However, atelectasis can also be asymptomatic. Fever is not a symptom of atelectasis.

Atelectasis can be caused by various medical conditions, including post-surgical complications, surfactant deficiency, and poor surfactant spreading during inspiration. It can also be caused by blockage of a bronchiole or bronchus, such as by a foreign body, mucus plug, tumor, or compression from the outside. Risk factors for atelectasis include certain types of surgery, muscle relaxation, obesity, high oxygen, lower lung segments, age, chronic obstructive pulmonary disease (COPD), asthma, and type of anesthetic.

The relevant symptoms for atelectasis include cough (not prominent), chest pain (not common), breathing difficulty (fast and shallow), low oxygen saturation, pleural effusion (transudate type), cyanosis (late sign), and increased heart rate.

The relevant clinical signs for the etiological diagnosis of atelectasis may include cough, chest pain (not common), breathing difficulty (fast and shallow), low oxygen saturation, pleural effusion (transudate type), cyanosis (late sign), and increased heart rate. However, it is important to note that atelectasis may also be asymptomatic.

The relevant laboratory data for the etiological diagnosis of atelectasis are not provided in the given information.

The relevant clinical characteristics for the etiological diagnosis of atelectasis include cough (not prominent), chest pain (not common), breathing difficulty (fast and shallow), low oxygen saturation, pleural effusion (transudate type), cyanosis (late sign), and increased heart rate. It is important to note that atelectasis does not cause fever. The underlying causes of atelectasis can include adjacent compression, passive atelectasis, dependent atelectasis, and poor surfactant spreading. Risk factors for atelectasis include type of surgery, use of muscle relaxation, obesity, high oxygen, lower lung segments, age, presence of chronic obstructive pulmonary disease or asthma, and type of anesthetic. Diagnosis of atelectasis is generally confirmed through chest X-ray, which may show lung opacification and/or loss of lung volume. Additional imaging modalities such as CT chest or bronchoscopy may be necessary to determine the cause of atelectasis.

The patient's personal relevant history for the etiological diagnosis of atelectasis includes post-surgical atelectasis as a common cause, as well as pulmonary tuberculosis, smoking, and old age as risk factors. Other factors associated with the development of atelectasis include the presence of chronic obstructive pulmonary disease or asthma, and the type of anesthesia used. The diagnosis of atelectasis is generally confirmed through chest X-ray, which shows small volume linear shadows, usually peripherally or at the lung bases. CT chest or bronchoscopy may be necessary to determine the cause or confirm the absence of proximal obstruction .\\

=====
Make a prediction based on your prior knowledge and on the patient's report. Please check if your prediction is in accordance with your prior knowledge.

Ensure your answers are using following template.

1. anesthesia: {True / False}, because...
2. asthma: {True / False}, because...\\
3. asymptomatic: {True / False}, because...
4. breathing difficulty: {True / False}, because...
5. chest pain: {True\textbackslash False}, because...
6. chronic obstructive pulmonary disease: {True / False}, because...
7. cough: {True / False}, because...
8. cyanosis: {True / False}, because...
9. fever: {True / False}, because...
10. pleural effusion: {True / False}, because...
11. pulmonary tuberculosis: {True / False}, because...
12. small volume linear shadows: {True / False}, because...
13. smoking: {True\textbackslash False}, because...
14. oxygen saturation: {True\textbackslash False}, because...
- ...

Fig. A.7. Concrete example of the prompt used in the Synthetic Clinical Feature Value Generation Module after the Augmented Clinical Knowledge Module generated new features.

Table B.5
List of all features identified by DALL-M.

Augmented features		
Abdominal liquid	Fungal and mineral exposure	Pleural thickening
Abdominal pain on right superior quadrant	Gender	Pneumonia
Age	Heart failure	Pneumothorax
Anesthesia	Heart rate	Positional chest pain
Asymptomatic	Hiatal hernia	Pulmonary tuberculosis
Back pain	High bone density	Pulsus paradoxus
Bone fractures	High heart beat intensity	Radiation to the trapezius ridge
Breathing difficulty	Hypertensive crisis	Relief of pain by bending forward
Cardiac tamponade	Hypoxemia	Relief of pain by inspiration
Chest Pain	Impaired gas exchange	Respiratory failure
Chronic obstructive pulmonary disease	Impaired left ventricular function	Respiratory rate
Congestive heart failure	Increased jugular venous pressure	Shortness of breath
Cough	Increased microvascular permeability	Small volume linear shadows
Decreased alertness	Interstitial lung disease	Smoking
Diaphoresis	Left precordial pleuritic chest pain	Specific electrocardiogram changes
Diastolic blood pressure (mmHg)	Leg swelling	Spirometry (FVC)
Dimension of the jugular vein	Lip cyanosis	Sudden onset of sharp chest pain
Distant heart sounds	Low heart beat intensity	Sweating
Dry cough	Lung collapsed	Systolic blood pressure (mmHg)
Dyspnea	Neck pain	Temperature
Electrical alternans on ekg	Oxygen saturation (%)	Venous pressure
Fatigue	Pain in the shoulders	Volume overload
Fever	Palpitations	Weakness
Fluid around the heart	Pericardial effusion	Wheezing
Friction rub	Plaques	Worsening of pain by inspiration
Frothy sputum	Pleural effusion	Worsening of pain by lying down

References

- [1] L. Liu, X. Yang, J. Lei, X. Liu, Y. Shen, Z. Zhang, P. Wei, J. Gu, Z. Chu, Z. Qin, et al., A survey on medical large language models: Technology, application, trustworthiness, and future directions, 2024, arXiv preprint arXiv:2406.03712.
- [2] V. Borisov, K. Seßler, T. Leemann, M. Pawelczyk, G. Kasneci, Language models are realistic tabular data generators, 2022, arXiv preprint arXiv:2210.06280.
- [3] B. Theodorou, C. Xiao, J. Sun, Synthesize high-dimensional longitudinal electronic health records via hierarchical autoregressive language model, *Nat. Commun.* 14 (1) (2023) 5305.
- [4] A. Biseric, M. Nikolic, M. van der Schaar, B. Delibasic, P. Lio, A. Petrovic, Interpretable medical diagnostics with structured data extraction by large language models, 2023, arXiv preprint arXiv:2306.05052.
- [5] R.J. Chen, M.Y. Lu, T.Y. Chen, D.F. Williamson, F. Mahmood, Synthetic data in machine learning for medicine and healthcare, *Nat. Biomed. Eng.* 5 (6) (2021) 493–497.
- [6] A. Gonzales, G. Guruswamy, S.R. Smith, Synthetic data in health care: A narrative review, *PLOS Digit. Heal.* 2 (1) (2023) e0000082.
- [7] W. Ollier, T. Sprosen, T. Peakman, UK Biobank: From concept to reality, *Pharmacogenomics* 6 (6) (2005) 639–646.
- [8] L. Cui, H. Li, K. Chen, L. Shou, G. Chen, Tabular data augmentation for machine learning: Progress and prospects of embracing generative AI, 2024, arXiv preprint arXiv:2407.21523.
- [9] G. Fan, J. Wang, Y. Li, D. Zhang, R.J. Miller, Semantics-aware dataset discovery from data lakes with contextualized column-based representation learning, *Proc. VLDB Endow.* 16 (7) (2023) 1726–1739.
- [10] X. Fang, W. Xu, F. Tan, J. Zhang, Z. Hu, Y. Qi, S. Nickleach, D. Socolinsky, S. Sengamedu, C. Faloutsos, Large language models (LLMs) on tabular data: Prediction, generation, and understanding—A survey, 2024, arXiv 2024. arXiv preprint arXiv:2402.17944.
- [11] P. Li, Y. He, D. Yashar, W. Cui, S. Ge, H. Zhang, D. Rifinski Fainman, D. Zhang, S. Chaudhuri, Table-GPT: Table fine-tuned GPT for diverse table tasks, *Proc. ACM Manag. Data* 2 (3) (2024) 1–28.
- [12] X. Tang, Y. Zong, J. Phang, Y. Zhao, W. Zhou, A. Cohan, M. Gerstein, Struc-bench: Are large language models really good at generating complex structured data?, 2023, arXiv preprint arXiv:2309.08963.
- [13] A. Nazabal, P.M. Olmos, Z. Ghahramani, I. Valera, Handling incomplete heterogeneous data using vaes, *Pattern Recognit.* 107 (2020) 107501.
- [14] J. Engelmann, S. Lessmann, Conditional Wasserstein GAN-based oversampling of tabular data for imbalanced learning, *Expert Syst. Appl.* 174 (2021) 114582.
- [15] T. Liu, J. Fan, N. Tang, G. Li, X. Du, Controllable tabular data synthesis using diffusion models, *Proc. ACM Manag. Data* 2 (1) (2024) 1–29.
- [16] M. Yakout, K. Ganjam, K. Chakrabarti, S. Chaudhuri, Infogather: Entity augmentation and attribute discovery by holistic matching with web tables, in: *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, 2012, pp. 97–108.
- [17] S. Zhang, K. Balog, Web table extraction, retrieval and augmentation, in: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 1409–1410.
- [18] F. Nargesian, E. Zhu, K.Q. Pu, R.J. Miller, Table union search on open data, *Proc. VLDB Endow.* 11 (7) (2018) 813–825.
- [19] A. Bogatu, A.A. Fernandes, N.W. Paton, N. Konstantinou, Dataset discovery in data lakes, in: *2020 IEEE 36th International Conference on Data Engineering, ICDE, IEEE*, 2020, pp. 709–720.
- [20] A.D. Sarma, L. Fang, N. Gupta, A.Y. Halevy, H. Lee, F. Wu, R. Xin, C. Yu, Finding related tables, in: *SIGMOD Conference*, vol. 10, 2012, pp. 2213836–2213962.
- [21] G. Limaye, S. Sarawagi, S. Chakrabarti, Annotating and searching web tables using entities, types and relationships, *Proc. VLDB Endow.* 3 (1–2) (2010) 1338–1347.
- [22] S. Zhang, K. Balog, Entitables: Smart assistance for entity-focused tables, in: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2017.
- [23] L. Zhang, S. Zhang, K. Balog, Table2Vec: Neural word and entity embeddings for table population and retrieval, in: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019.
- [24] A. Khatiwada, G. Fan, R. Shraga, Z. Chen, W. Gatterbauer, R.J. Miller, M. Riedewald, Santos: Relationship-based semantic table union search, *Proc. the ACM Manag. Data* 1 (1) (2023) 1–25.
- [25] M. Yakout, K. Ganjam, K. Chakrabarti, S. Chaudhuri, Infogather: Entity augmentation and attribute discovery by holistic matching with web tables, in: *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, 2012, pp. 97–108.
- [26] P. Chen, S. Sarkar, L. Lausen, B. Srinivasan, S. Zha, R. Huang, G. Karypis, Hytrel: Hypergraph-enhanced tabular data representation learning, *Adv. Neural Inf. Process. Syst.* 36 (2024).
- [27] R. Cappuzzo, P. Papotti, S. Thirumuruganathan, Creating embeddings of heterogeneous relational datasets for data integration tasks, in: *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, 2020, pp. 1335–1349.
- [28] X. Hu, S. Wang, X. Qin, C. Lei, Z. Shen, C. Faloutsos, A. Katsifodimos, G. Karypis, L. Wen, S.Y. Philip, Automatic table union search with tabular representation learning, in: *Findings of the Association for Computational Linguistics, ACL 2023*, 2023, pp. 3786–3800.
- [29] G. Fan, J. Wang, Y. Li, D. Zhang, R. Miller, Semantics-aware dataset discovery from data lakes with contextualized column-based representation learning, 2022, arXiv preprint arXiv:2210.01922.
- [30] E. Zhu, D. Deng, F. Nargesian, R.J. Miller, Josie: Overlap set similarity search for finding joinable tables in data lakes, in: *Proceedings of the 2019 International Conference on Management of Data*, 2019, pp. 847–864.
- [31] E. Zhu, F. Nargesian, K.Q. Pu, R.J. Miller, LSH ensemble: Internet-scale domain search, *Proc. VLDB Endow.* 9 (12) (2016) 1185–1196.
- [32] N. Chepurko, R. Marcus, E. Zraggen, R.C. Fernandez, T. Kraska, D. Karger, Arda: Automatic relational data augmentation for machine learning, 2020, arXiv preprint arXiv:2003.09758.
- [33] Y. Dai, A. El-Roby, E. Adee, V. Thaker, OmniMatch: Overcoming the cold-start problem in cross-domain recommendations using auxiliary reviews, in: *Proceedings of the 28th International Conference on Extending Database Technology, EDBT*, 2025, pp. 80–91.

- [34] Y. Dong, K. Takeoka, C. Xiao, M. Oyamada, Efficient joinable table discovery in data lakes: A high-dimensional similarity-based approach, in: 2021 IEEE 37th International Conference on Data Engineering, ICDE, IEEE, 2021, pp. 456–467.
- [35] Y. Dong, C. Xiao, T. Nozawa, M. Enomoto, M. Oyamada, Deepjoin: Joinable table discovery with pre-trained language models, 10, 2022, p. 2022, URL: <https://arxiv.org/abs/2212.07588>.
- [36] J. Liang, C. Lei, X. Qin, J. Zhang, A. Katsifodimos, C. Faloutsos, H. Rangwala, Featnavigator: Automatic feature augmentation on tabular data, 2024, arXiv preprint arXiv:2406.09534.
- [37] S. Bharadwaj, P. Gupta, R. Bhagwan, S. Guha, Discovering related data at scale, *Proc. VLDB Endow.* 14 (8) (2021) 1392–1400.
- [38] Z. Zhao, R. Castro Fernandez, Leva: Boosting machine learning performance with relational embedding data augmentation, in: Proceedings of the 2022 International Conference on Management of Data, 2022, pp. 1504–1517.
- [39] M. Glass, X. Wu, A.R. Naik, G. Rossiello, A. Gliozzo, Retrieval-based transformer for table augmentation, 2023, arXiv preprint arXiv:2306.11843.
- [40] Y. Gu, R. Tinn, H. Cheng, M. Lucas, N. Usuyama, X. Liu, T. Naumann, J. Gao, H. Poon, Domain-specific language model pretraining for biomedical natural language processing, *ACM Trans. Comput. Heal.* 3 (1) (2021) 1–23.
- [41] T. Han, L.C. Adams, J.M. Papaioannou, P. Grundmann, T. Oberhauser, A. Löser, D. Truhn, K.K. Bressen, Medalpaca—An open-source collection of medical conversational AI models and training data, 2023, arXiv preprint arXiv:2304.08247.
- [42] A. Khatiwada, R. Shraga, W. Gatterbauer, R.J. Miller, Integrating data lake tables, *Proc. VLDB Endow.* 16 (4) (2022) 932–945.
- [43] R.J. Miller, Open data integration, *Proc. VLDB Endow.* 11 (12) (2018) 2130–2139.
- [44] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, K. Talwar, Privacy, accuracy, and consistency too: A holistic solution to contingency table release, in: Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, 2007, pp. 273–282.
- [45] J. Zhang, G. Cormode, C.M. Procopiuc, D. Srivastava, X. Xiao, Privbayes: Private data release via Bayesian networks, *ACM Trans. Database Syst.* 42 (4) (2017) 1–41.
- [46] N. Park, M. Mohammadi, K. Gorde, S. Jajodia, H. Park, Y. Kim, Data synthesis based on generative adversarial networks, *Proc. VLDB Endow.* 11 (10) (2018) 1071–1083.
- [47] J. Jordon, J. Yoon, M. Van Der Schaar, Pate-GAN: Generating synthetic data with differential privacy guarantees, in: International Conference on Learning Representations, 2018.
- [48] H. Chen, S. Jajodia, J. Liu, N. Park, V. Sokolov, V. Subrahmanian, Faketables: Using gans to generate functional dependency preserving tables with bounded real data, in: Proceedings of the International Joint Conference on Artificial Intelligence, 2019, pp. 2074–2080.
- [49] Y. Zhang, N. Zaidi, J. Zhou, G. Li, Interpretable tabular data generation, *Knowl. Inf. Syst.* 65 (7) (2023) 2935–2963.
- [50] R. Al-Bahrani, D. Jha, Q. Kang, S. Lee, Z. Yang, W.K. Liao, A. Agrawal, A.N. Choudhary, SIGRNN: Synthetic minority instances generation in imbalanced datasets using a recurrent neural network, in: Proceedings of the International Conference on Pattern Recognition Applications and Methods, 2021, pp. 349–356.
- [51] J. Engelman, S. Lessmann, Conditional Wasserstein GAN-based oversampling of tabular data for imbalanced learning, *Expert Syst. Appl.* 174 (2021) 114582.
- [52] J. Kim, C. Lee, Y. Shin, S. Park, M. Kim, N. Park, J. Cho, SOS: Score-based oversampling for tabular data, in: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 762–772.
- [53] L. Xu, M. Skoularidou, A. Cuesta-Infante, K. Veeramachaneni, Modeling tabular data using conditional gan, in: Advances in Neural Information Processing Systems, *NeurIPS*, vol. 32, 2019.
- [54] Y. Lin, B. Ding, H. Jagadish, J. Zhou, Smartfeat: Efficient feature construction through feature-level foundation model interactions, 2023, arXiv preprint arXiv:2309.07856.
- [55] J.M. Kanter, K. Veeramachaneni, Deep feature synthesis: Towards automating data science endeavors, in: 2015 IEEE International Conference on Data Science and Advanced Analytics, DSAA, IEEE, 2015, pp. 1–10.
- [56] G. Katz, E.C.R. Shin, D. Song, Explores: Automatic feature generation and selection, in: 2016 IEEE 16th International Conference on Data Mining, ICDM, IEEE, 2016, pp. 979–984.
- [57] D. Wang, M. Xiao, M. Wu, P. Wang, Y. Zhou, Y. Fu, Reinforcement-enhanced autoregressive feature transformation: Gradient-steered search in continuous space for postfix expressions, in: Advances in Neural Information Processing Systems, vol. 36, 2023.
- [58] M. Xiao, D. Wang, M. Wu, P. Wang, Y. Zhou, Y. Fu, Beyond discrete selection: Continuous embedding space optimization for generative feature selection, in: 2023 IEEE International Conference on Data Mining, ICDM, 2023, pp. 688–697.
- [59] S. Van Buuren, K. Groothuis-Oudshoorn, Mice: Multivariate imputation by chained equations in R, *J. Stat. Softw.* 45 (2011) 1–67.
- [60] D.J. Stekhoven, P. Bühlmann, MissForest—non-parametric missing value imputation for mixed-type data, *Bioinformatics* 28 (1) (2012) 112–118.
- [61] Z. Wang, J. Sun, Transtab: Learning transferable tabular transformers across tables, in: Advances in Neural Information Processing Systems, vol. 35, 2022, pp. 2902–2915.
- [62] Y. Tian, R. Gan, Y. Song, J. Zhang, Y. Zhang, Chimed-GPT: A Chinese medical large language model with full training regime and better alignment to human preferences, 2023, arXiv preprint arXiv:2311.06025.
- [63] M. Xu, MedicalGPT: Training medical gpt model, 2023, <https://github.com/shibing624/MedicalGPT>.
- [64] J. Chen, X. Wang, A. Gao, F. Jiang, S. Chen, H. Zhang, D. Song, W. Xie, C. Kong, J. Li, et al., HuatuoGPT-II, one-stage training for medical adaptation of LLMs, 2023, arXiv preprint arXiv:2311.09774.
- [65] W. Zhu, X. Wang, ChatMed: A Chinese medical large language model, 2023.
- [66] Q. Lu, D. Dou, T. Nguyen, Clinical5: A generative language model for clinical text, in: Findings of the Association for Computational Linguistics, EMNLP 2022, 2022, pp. 5436–5443.
- [67] G. Wang, G. Yang, Z. Du, L. Fan, X. Li, ClinicalGPT: Large language models finetuned with diverse medical data and comprehensive evaluation, 2023, arXiv preprint arXiv:2306.09968.
- [68] R. Luo, L. Sun, Y. Xia, T. Qin, S. Zhang, H. Poon, T.Y. Liu, BioGPT: Generative pre-trained transformer for biomedical text generation and mining, *Brief. Bioinform.* 23 (6) (2022) <http://dx.doi.org/10.1093/bib/bbac409>.
- [69] X. Yang, A. Chen, N. Pournejatian, H.C. Shin, K.E. Smith, C. Parisien, C. Compas, C. Martin, A.B. Costa, M.G. Flores, et al., A large language model for electronic health records, *NPJ Digit. Med.* 5 (1) (2022) 194.
- [70] K. Singhal, S. Azizi, T. Tu, S.S. Mahdavi, J. Wei, H.W. Chung, N. Scales, A. Tanwani, H. Cole-Lewis, S. Pfohl, et al., Large language models encode clinical knowledge, *Nature* 620 (7972) (2023) 172–180.
- [71] C. Li, C. Wong, S. Zhang, N. Usuyama, H. Liu, J. Yang, T. Naumann, H. Poon, J. Gao, LLaVA-Med: Training a large language-and-vision assistant for biomedicine in one day, *Adv. Neural Inf. Process. Syst.* 36 (2024).
- [72] R. Bigolin Lanfredi, M. Zhang, W.F. Auffermann, J. Chan, P.A.T. Duong, V. Srikumar, T. Drew, J.D. Schroeder, T. Tasdizen, Reflax, a dataset of reports and eye-tracking data for localization of abnormalities in chest X-rays, *Sci. Data* 9 (1) (2022) 350.
- [73] Radiopedia, 2023, <https://radiopaedia.org/>.
- [74] Wikipedia, 2023, <https://wikipedia.org/>.
- [75] A. Johnson, L. Bulgarelli, T. Pollard, S. Horng, L.A. Celi, R. Mark, Mimic-IV, 2021.
- [76] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpanskaya, et al., Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 590–597.
- [77] K. Huang, J. Altsaar, R. Ranganath, Clinicalbert: Modeling clinical notes and predicting hospital readmission, 2019, <http://dx.doi.org/10.48550/ARXIV.1904.05342>, <https://arxiv.org/abs/1904.05342>.
- [78] E. Alsentzer, J.R. Murphy, W. Boag, W.H. Weng, D. Jin, T. Naumann, M.B.A. McDermott, Publicly available clinical bert embeddings, 2019.
- [79] Z. Chen, A.H. Cano, A. Romanou, A. Bonnet, K. Matoba, F. Salvi, M. Pagliardini, S. Fan, A. Köpf, A. Mohtashami, A. Sallinen, A. Sakhaeirad, V. Swamy, I. Krawczuk, D. Bayazit, A. Marmet, S. Montariol, M.A. Hartley, M. Jaggi, A. Bosselut, Meditron-70b: Scaling medical pretraining for large language models, 2023.
- [80] A.Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D.S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L.R. Lavaud, M.A. Lachaux, P. Stock, T.L. Scao, T. Lavril, T. Wang, T. Lacroix, W.E. Sayed, Mistral 7b, 2023.
- [81] L. Tunstall, E. Beeching, N. Lambert, N. Rajani, K. Rasul, Y. Belkada, S. Huang, L.von Werra, C. Fourrier, N. Habib, N. Sarrazin, O. Sanseviero, A.M. Rush, T. Wolf, Zephyr: Direct distillation of LM alignment, 2023.
- [82] Hugo Touvron, E.A. Llama 2: Open foundation and fine-tuned chat models, 2023.
- [83] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16 × 16 words: Transformers for image recognition at scale, 2020, arXiv preprint arXiv:2010.11929.
- [84] Y. Gorishniy, I. Rubachev, V. Khurlov, A. Babenko, Revisiting deep learning models for tabular data, *Adv. Neural Inf. Process. Syst.* 34 (2021) 18932–18943.
- [85] L. Grinsztajn, E. Oyallon, G. Varoquaux, Why do tree-based models still outperform deep learning on typical tabular data?, 35, 2022, pp. 507–520.
- [86] C. Peng, X. Yang, A. Chen, et al., A study of generative large language model for medical research and healthcare, *NPJ Digit. Med.* 6 (1) (2023) 210.
- [87] A.J. Thirunavukarasu, D.S.J. Ting, K. Elangovan, L. Gutierrez, T.F. Tan, D.S.W. Ting, Large language models in medicine, *Nature Med.* 29 (8) (2023) 1930–1940.
- [88] Y.L. Chou, C. Moreira, B. Bruza, C. Ouyang, J. Jorge, Counterfactuals and causability in explainable artificial intelligence: Theory, algorithms, and applications, *Inf. Fusion* 81 (2022) 59–83.