

A Work Project, presented as part of the requirements for the Award of a Master's degree in
Business Analytics from the Nova School of Business and Economics.

**Predictive Modeling for Clinical Trial Completion: Assessing the Phase Success -
Incorporating RAG Techniques for Predictive Modeling of Clinical Trial Outcomes**

Jasmin Hamrouni

Work project carried out under the supervision of:

Qiwei Han

17/12/2024

Abstract (100 words maximum)

This study investigates predictive modeling of clinical trial completion using the *HINTBasic* and *HINTPlus* models. By integrating multimodal datasets, the models predict clinical trial phase success. It provides interpretability insights into the *HINTPlus* model's decision-making process. Retrieval-Augmented-Generation techniques were used to contextualize results. Our findings support informed decision-making, optimize resource allocation, and accelerate drug development in clinical trials.

Keywords (Clinical Trials, Health Care, Artificial Intelligence, Machine Learning Methods, Predictive Modeling, , Model Interpretability, Contextuality Retrieval-Augmented-Generation)

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22

Index

1. Introduction.....	4
1.1 Recent Industry Insights.....	6
1.2 Limits of Clinical Trials.....	7
1.3 Literature Review and Limitations of Existing Models.....	8
1.4 Proposed Model: Hierarchical Interaction Network (<i>HINT</i>).....	11
2. Literature review.....	13
3. Exploratory Data Analysis.....	21
3.1 Data Collection and Description.....	22
3.2 Data Cleaning and Preprocessing.....	22
3.3 Insights and Patterns.....	24
3.3.1 Study Completion.....	24
3.3.2 Study Duration.....	26
3.3.3 Study design.....	26
3.3.4 Primary purposes of clinical trials.....	27
3.3.5 Patient Characteristics.....	28
3.3.6 Enrollment.....	29
3.3.7 Contribution.....	29
4. Benchmark.....	30
4.1 Disease data.....	31
4.2 Drug molecule data.....	32
4.3 Trial eligibility criteria.....	33
4.4 Trial outcome information.....	33
4.5 <i>ADMET</i>	33
4.6 Enrollment.....	34
4.7 Data preprocessing.....	35
4.8 Data split.....	36
5. <i>HINTBasic</i> model.....	38
6. <i>HINTPlus</i> Model.....	43
7. Results.....	44
7.1 Evaluation of Our Models.....	46
7.2 Comparison between <i>HINTBasic</i> and <i>HINTPlus</i>	47
7.3 Evaluation of the <i>HINTBasic</i> with baseline models.....	47
7.3.1 Phase I.....	48

7.3.2	Phase II.....	48
7.3.3	Phase III.....	48
7.4	Disease groups evaluation.....	49
7.5	Results overview.....	50
8.	Discussion.....	51
8.1	Data-related limitations.....	52
8.2	Cost/Benefit analysis.....	53
8.2.1	Benefits.....	53
8.2.2	Lower costs implications.....	54
8.3	Business implications based on stakeholders.....	55
9.	Conclusion.....	58
10.	Incorporating RAG Techniques for Predictive Modeling of Clinical Trial Outcomes	61
10.1	Introduction.....	61
10.2	Literature Review.....	62
10.3	Methodology.....	64
10.3.1	Database.....	64
10.3.2	RAG Architectural Structure.....	64
10.3.3	Experimental Setup and Hypotheses.....	66
10.4	Results.....	67
10.4.1	Retrieval.....	67
10.4.2	Generation.....	68
10.5	Discussion.....	69
10.5.1	Architectural Limitations.....	69
10.5.2	Stakeholder Implementation.....	70
10.6	Conclusion.....	71

1. Introduction

20 million new cases and 9.7 million deaths. These are the numbers reported by the World Health Organization (WHO)'s cancer agency (2024), describing the state of the growing worldwide burden of cancer in 2022. Their analysis also shows that 1 in 5 people develop cancer in the span of their lives. A look into the future suggests that over 35 million new cases are predicted in 2050, which is a 77% increase in comparison to the estimated 20 million cases in 2022. According to Dr. Adams, the Head of the Union for International Cancer Control, governments need to facilitate cancer care, so that all humans have access. Yet he also highlights the progress which has already been made in early detection of cancer as well as their treatment. One way to access the newest, most innovative treatments is through participation in clinical trials. For example, Singh et al. (2016) demonstrated the impact of clinical trials in cancer treatment, a type of blood cancer called acute lymphoblastic leukemia, which led to remission rates of 90%. To understand clinical trials in a broader way, we will now dive deeper into their inner workings, as well as their impact on the market.

Clinical trials are a multi-phase process designed to assess the safety and efficacy of new drugs or treatment modalities before the official approval by the authorities for public use (Hay et al, 2014).

Clinical trials are an essential element of the drug development process, and their importance cannot be overstated, as they stand as the gold standard in medical research (Ghim & Ahn 2023). Despite their critical role, the process of conducting a trial carries a long list of inherent complexities, high costs, and time-consuming nature.

Indeed, the strict regulatory compliance (dictated by regulatory bodies such as the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA)) imposes rigorous standards to ensure both safety and reliability. The regulatory landscape additionally involves the adherence to Good Clinical Practice (GCP) guidelines and the observance of

ethical considerations such as informed consent, patient privacy, and the use of placebos (National Institute of Health, World Health Organization).

Clinical trials are typically divided into Phases I, II, III, and IV. Phase I trials primarily focus on determining the safety of a new drug and identifying potential side effects, often involving a small group of healthy volunteers or patients. Around 20-80 patients are involved in this phase (National Institute of Health). However, only 70 % of these patients will progress to the next stage in 3-6 months. Phase II trials expand the focus to evaluate the efficacy of the drug and further assess its safety, usually involving a larger group of patients with the target condition. Phase III trials aim to confirm the drug's efficacy, monitor side effects, and compare it to existing standard treatments, involving a much larger patient population to ensure statistically significant results (Chopra et al 2024, 4212). In fact, Phase III is subject to 10 times more participants in comparison to Phase II, to establish efficacy, monitor adverse effects as well as compare the results to other treatments (National Institute of Health; Getz et al. 2016). This process has a duration of 1 to 4 years and only one fourth of participants will progress to the next phase (Chopra et al. 2024, 4212).

Phase IV trials are completed trials that study the side effects caused over time by a new treatment after it has been approved and is on the market. These could also be referred to as post marketing surveillance trials. The aim is to look for side effects that have not been detected in earlier phases of trials and may also assess the performance of a new treatment over a long period of time (National Cancer Institute). This phase might take a year or more to come to an end and it has a success rate of 70-90% (Chopra et al. 2024, 4212).




Figure 1 - Clinical Trials Cost distribution (Clinical Research IO 2018)

This essential process is also one of the most resource-intensive components of drug development in terms of both financial and time perspective. Indeed, the average time to bring a new drug to market exceeds a decade, while the total development is estimated to be \$2.6 billion, with clinical trials accounting for a major part of these expense (DiMasi et al. 2016).

1.1 Recent Industry Insights

With regards to the market trends, it can be stated that the clinical trials field is expected to be thriving in the next decade, and major investments are being placed to favor the development of AI practices aimed at achieving improved efficacy and reducing expenditure. Recent industry reports provide valuable insights related to the current state and future projections of the industry.

An industry report by the Global Market Insights from 2024 has shown that the global clinical trials market has been valued at approximately \$ 55.8 billion in 2023, with an annual growth rate (CAGR) of 5.4 %. This will lead to it reaching a value of around \$ 89.8 billion by 2032. (2024) have estimated that the global clinical trial market has been valued at \$ 57.76 billion in 2023, and it is expecting a CAGR of 7.1 %, therefore reaching \$ 106.78 billion by 2032. (Faizullabhoy et al. 2024)

A closer look at clinical trials suggests that Phase III trials dominate the market, accounting for 53.3 % of the total revenue share in 2023. Phase II trials are also significant, with the segment projected to lead the market, accounting for the largest revenue of \$ 23.4 billion in

2023 (Grand View Research 2024). A huge priority in treatments is oncology as it remains a primary focus, with a substantial number of trials initiated in this area. Other significant therapeutic areas include central nervous system disorders, cardiology, and infectious diseases (Grand View Research 2024).

The market also experiences a shifting focus when it comes to countries and regions in which these clinical trials are carried out. Europe's share of global commercial clinical drug trials shows a significant decrease, declining from 22 % in 2013 to 12 % in 2023. On the contrary, China has doubled its share, accounting for 18 % of the global total, while the US remains the leader despite a slight decrease. This might be explained by pharmaceutical companies favoring US and China for their more straightforward regulatory environments (Financial Times 2024).

1.2 Limits of Clinical Trials

While clinical trials are crucial for drug development, they are also limited by several constraints, primarily related to financial resources and time.

From an historical point of view, the cost of clinical trials has grown significantly over the past few decades. Specifically, it doubled every nine years since the 1950s, following a trend described by the Eroom's Law, the reverse of Moore's Law (Scannell et al. 2012).

Furthermore, Phase III clinical trials account for most of the expenditure, with cost spanning between \$11 million and \$52 million, depending on the therapeutic area and the study requirements (Mestre-Ferrandiz et al. 2012).

Another element that could enhance the overall financial burden is represented by the large-scale patient recruitment and the need for extensive data collection (Getz et al. 2016).

In terms of duration, clinical trials are notoriously lengthy. For instance, the average duration of each phase is approximately 6 or 7 years, even though later phases often exceed this estimation; hence, the total development process timeline exceeds the decade and could reach

15 years (Wouters et al. 2020). Likewise, the process of patient recruitment is another major challenge. Finding suitable candidates who meet all the eligibility criteria is extremely time-consuming and maintaining participant retention throughout the duration of a trial can be equally difficult. The process of patient recruitment often involves identifying individuals who meet strict inclusion and exclusion criteria, which may include factors such as age, gender, medical history, and current health status. This can be further hindered by a lack of awareness or understanding of clinical trials among potential participants, as well as logistical challenges such as travel requirements, and the time commitment involved in participating in a trial (Bieganeck et al. 2022). This is enhanced in the case of trials focused on rare disease or demographic-specific studies, where the pool of eligible participants is already limited. Moreover, it is to be emphasized that the relevant failure rate of clinical trials. Failure root causes may include a broad variety of diverse factors, which range from the lack of efficacy or unforeseen safety issues to difficulties in managing the complexities of trial design (Getz et al. 2016; DiMasi et al. 2016).

1.3 Literature Review and Limitations of Existing Models

Artificial Intelligence has the power to be a real shift in the medical field as it could be able to accelerate the efficiency, economy and timeliness of drug development. In recent years, there have been notable advancements in artificial intelligence (AI) and machine learning, which have brought innovative techniques to the field of clinical trials, offering solutions to overcome challenges (Chopra et al. 2024).

AI-driven models have proven their ability to optimize various aspects of trials processes, such as patient recruitment, outcome prediction, and synthetic data generation. Specifically, the integration of Large Language Models (LLMs) has meaningfully facilitated the process by analyzing large datasets, thereby improving trial design and forecasting outcomes (McKinsey 2024).

The models developed in latest years, such as those developed by Kavalci & Hartshorn (2023), have delivered promising methodologies for predicting trial success and preventing early terminations. AI practices also suggested hypothetical resolutions to address the challenges of patient recruitment and retention. Indeed, through the examination of electronic health records (EHRs) and other real-world data, AI models can be leveraged to identify participants who fulfil eligibility criteria for a given trial, thus streamlining the recruitment process. Their predictive ability could also be employed to forecast patient drop-out rates and detect factors that may influence retention at an early stage, therefore letting trial organizers take proactive measures to preserve high engagement. Similarly, AI-driven approaches have been hired with the purpose of enhancing trial efficiency, by simulating diverse scenarios in which parameters such as dosage, sample size, and endpoint selection have been varied accordingly. These techniques evidenced their capability to shorten trial timelines and improve the overall efficiency of the drug development process (Lu et al. 2024). According to a study conducted by McKinsey in November 2023, the integration of AI into clinical studies might potentially reduce trials timeline by 15-30%.

Despite the enormous advancements, the models proposed thus far also have several limitations. A primary obstacle lies in the reliance on large amounts of high-quality data, which is often not available. For instance, Large Language Models, while highly effective, require extensive labelled datasets to accomplish accurate predictions, hence their applicability is only feasible in areas where data is consistent and reliable. Similarly, as many of the existing models are based on binary outcomes (trials are classified as either successful or unsuccessful) the complexity of clinical research is not fully captured, resulting in a more nuanced interpretation that does not underscore determinant factors for success or failure (Topol 2019; Yu et al. 2018).

Another limitation of existing AI models might be found in the lack of transparency and interpretability, since numerous AI-driven models, especially those relying on deep learning techniques, act as hidden frameworks and do not clearly show how predictions are made. This might pose as a barrier, blocking these models from being adopted in clinical trials, as regulatory bodies and stakeholders require a clearer understanding of the factors affecting trial outcomes. Still, the models currently in use recurrently fail to achieve generalizability, thereby their performance firmly depends on the context (Topol 2019; Yu et al. 2018).

According to Chopra et al. (2024), understanding and validating seems to be demanding, guidelines and standards for AI-driven healthcare solutions are being developed by authorities such as the FDA. However, this takes a huge amount of time, since this requires a long testing and validation phase, while also requiring explanation of the underlying algorithms. Deep learning neural networks are often called “black box” models because of their high level of difficulty, making it harder to understand. This complexity as well as its unclear interpretability of its result can also be a reason why mainstream adoption slows down. As a result, it might also be more challenging to get patients permission to use their data, due to the lack of understanding of what the data is being used for. Consequently, it is vital that AI models are transparent and easy to understand. Upholding a high patients confidence and a high standard of ethical behavior also holds when it comes to data privacy and data ownership to avoid unjust treatment. This might be the case when trials are using face recognition software or other methods to track if participants are sticking to the rules of the study. For that reason, it is important that researchers using AI and ML algorithms provide their participants with an explicit explanation of the risks and advantages of their data gathering (Chopra et al. 2024, 4212-4216).

Another closely related limitation is the inherent bias in the data used. Trained Algorithms on a specific dataset are at risk of excluding large parts of the populations which have not been

included. This is due to a serious lack of diversity in medical research, which can be traced back to research only including able bodied white people as the default. As a result, an AI trained model would not be able to have enough knowledge of underrepresented groups such as people of color and patients from lower socio-economic backgrounds, leading to biased findings that do not translate to the not included groups. One way to avoid this happening is if health equity consideration are implemented in AI and ML systems. Another way to treat this problem is the inclusion of domain experts such as medical professionals, which are consulted in the algorithm development process to help fill in missing values which are vital for providing context to the dataset (Chopra et al 2024, p.4218).

The limitations of traditional clinical trials, combined with the partial resolutions offered by existing AI models, emphasize the need for more sophisticated methodologies that can address these challenges comprehensively. There is a clear demand for multi-function approaches that could serve different purposes, such as integrating diverse data sources, providing interpretable insights, and generalizing adaptability across different trial settings, thereby proposing a more holistic solution to the issue faced by clinical trials.

1.4 Proposed Model: Hierarchical Interaction Network (*HINT*)

This thesis aims to address gaps identified in the current literature by implementing and furtherly refining the Hierarchical Interaction Network (*HINT*) model by Fu et al. (2022), which has demonstrated a strong potential in the prediction of clinical trial outcomes. *HINT* represents a novel approach that is characterized by the integration of a diverse dataset, including drug molecular structures, disease-specific information, trial eligibility criteria and pharmacokinetic interactions, to accurately predict the success or failure of clinical trials outcomes across different phases. Unlike traditional models, *HINT* leverages a hierarchical interaction graph to display the complex relationships among these components, providing a better comprehensive perspective of the influential factors.

Fu et al.'s (2022) methodology employed comprises several key components. First, multimodal data embedding is used to encode different types of trial-related information. The embeddings obtained are then combined with external knowledge sources, including pharmacokinetic data, to reach a deeper understanding of drug interactions and their effects on trial outcomes. To advance the progress of ML models in the context of clinical trials, *HINT* introduces TOP (Trial Outcome Prediction), a benchmark dataset, containing all the data sources used in the development of the model, providing a comprehensive resource for training and evaluation purposes.

The objective of this thesis is to develop an effective model for predicting the factors that influence the success or failure of clinical trials by leveraging *HINT*. We are defining success as the clinical trial reaching the end of its phase. This is done by collecting data, that has been used by Fu et al. (2022) for their benchmark, to recreate an updated and enriched version of TOP, which will then be leveraged to replicate the *HINT* model.

The results are expected to provide valuable insights into how AI-driven models can be used to improve clinical trial efficiency, reduce costs, and ultimately accelerate the approval of new therapies. By providing accurate predictions of trial outcomes, *HINT* could help pharmaceutical companies make more informed decisions about which drug candidates to advance, thereby reducing the risk of costly late-phase failures. Furthermore, the ability to predict trial outcomes at an early stage could enable more efficient allocation of resources, such as patient recruitment efforts and financial investments, ultimately making the entire drug development process more cost-effective.

Our work is structured as follows: firstly, a literature review provides a deep dive into the *HINT* model, which is followed by an evaluation of existing research, methodology and their limitations, addressing related research questions. Secondly, the data collection and preparation encoding process is explained, which is followed by the performance of an

Exploratory Data Analysis (EDA) containing insights as well as patterns and trends. This is followed by the model's learning phase. After, we explain the architecture of our model. In our results section, we present and evaluate our model's performance while also comparing it with other algorithms. Finally, we discuss the implications and limitations of our results and conclude key insights derived from our analysis, while also providing recommendations for future research.

In conclusion, this thesis will explore the potential of the *HINT* model to address some of the most pressing challenges in the field of clinical trials. The successful implementation of this model could pave the way for more widespread adoption of AI-driven approaches in drug development, ultimately benefiting patients by bringing new treatments to market faster and more efficiently. The implications of this research extend beyond the pharmaceutical industry, offering a framework that could be adapted to other areas of healthcare and medical research where complex, data-driven decision-making is required.

2. Literature review

Fu et al.'s (2022) study has a vital role in the field of clinical trial outcome predictions. By implementing the novel approach of the Hierarchical Interaction Network (*HINT*), which depicts the complex relationship between trial variables, predictive accuracy for clinical outcomes across phases is amplified. By capturing complex relationships among trial components, *HINT* can accurately predict the success or failure of clinical trials across different phases. Its ability to incorporate diverse data types and leverage external knowledge makes it one of the most comprehensive and effective tools in this domain. *HINT* functions as an end-to-end framework, which ultimately returns a success probability score for a clinical trial before it starts. By providing this insight, *HINT* could become a powerful tool for stakeholders to allocate resources in a more efficient way, while accelerating the clinical trials approval process.

Clinical trials are designed to assess the safety and efficacy of new treatments. A trial typically involves testing a treatment set, which includes one or more drug candidates, against a target disease set in a group of patients defined by eligibility criteria. The goal is to determine whether the treatment successfully meets its primary endpoints, such as reducing disease symptoms or achieving specific outcomes (Fu et al. 2022, 3-4).

In the *HINT* model, the outcome of a trial is indicated using a binary label: “1” indicates success, meaning that the trial met its primary endpoints, while “0” indicates failure.

HINT focuses on two key prediction tasks: the first is a phase-level prediction, which assesses the likelihood of success for a specific trial phase (e.g., Phase I, II, or III), while the second is an indication-level prediction, which judges whether a treatment will ultimately pass all three phases of the trial process (Fu et al. 2022, 4).

A carefully designed architecture that integrates multimodal data and external knowledge with modelling techniques achieves *HINT* predictive ability. The framework consists of several key components, featuring an input embedding module, a knowledge embedding module, a hierarchical interaction graph, and a dynamic attentive graph neural network (Fu et al. 2022, 4).

The first step consists in encoding multimodal data into embeddings, by using the input embedding module, which processes three primary data types: drug molecules, disease information, and trial protocols. Drug molecules are encoded through the SMILES (Simplified Molecular Input Line Entry System) strings and molecular graphs, which capture the structural and chemical properties of the drugs. To generate embeddings, these representations are processed through techniques such as Morgan fingerprints, SMILES encoders, and message-passing neural networks. The encoding of disease information is achieved by using hierarchical medical ontologies, such as ICD-10 codes and textual descriptions. *HINT* then leverages the Graph-based Attention Model to create embeddings

that reflect the hierarchical relationships inherent in these datasets. Trial protocols (including inclusion and exclusion criteria) are then encoded through Bio-BERT, a domain-specific language model able to detect the semantic notes of trial descriptions (Fu et al. 2022, 4-6).

Ultimately, the output is a set of embeddings that represent the key components of the trial. These embeddings are further enhanced by using external knowledge sources. Indeed, to provide a greater understanding of drugs interactions, the pharmacokinetics data present in *ADMET* dataset are incorporated (pharmacokinetics include properties such as absorption, distribution, metabolism, excretion, and toxicity).

Lastly, historical trial data, including success rates for specific diseases, are also integrated to inform predictions. These knowledge embeddings are pretrained on external datasets, enabling *HINT* to leverage vast amounts of prior information (Fu et al. 2022, 6-7). The following step consists in the hierarchical interaction graph, which lies at the heart of *HINT* architecture. This graph aims to connect the embeddings and establish relationships between drugs, diseases, and trial protocols. Therefore, it features several types of nodes, such as input nodes (representing the trial components), external knowledge nodes (representing *ADMET* properties and disease risks), aggregation nodes (which summarize interactions between trial components), and prediction nodes (which generate the final trial outcome predictions). This structure is essential to enable the model to capture both direct and indirect relationships, providing a holistic view of the trial's dynamics (Fu et al. 2022, 8-9).

After processing the graph, *HINT* aggregates the information from the graph to make the final prediction. The prediction node summarizes the information gathered from the pharmacokinetics properties, disease risk, and the interactions between the drug, disease, and eligibility criteria. The output of this prediction node is the final predicted trial outcome, which is typically a binary success/failure label, though it can also provide probability scores indicating the likelihood of success (Fu et al 2022, 8).

To refine predictions, an attentive graph neural network is employed. This component uses graph convolutional layers to aggregate information from neighboring nodes, updating the embeddings to reflect their context within the graph. Consequently, attention mechanisms assign weights to the most critical interactions, emphasizing the relationships that are most likely to influence trial outcomes, to improve predictive accuracy (Fu et al 2022, 8-9).




Figure 2 -HINT Framework, (Fu et al., 2022, Patterns 3, 100445) April 8, 2022, a 2022 The author(s)).

HINT was evaluated on the Trial Outcome Prediction (TOP) dataset, a comprehensive benchmark created to standardize clinical trial outcome prediction. The dataset includes 17,538 clinical trials, with 9,999 classified as successful and 7,539 as failures. For each trial, the dataset provides information on drug molecular structures, disease descriptions, trial protocols, and outcomes. Additional supplementary datasets include pharmacokinetics data (e.g., *ADMET* properties) and historical trial success rates. The TOP dataset spans a wide range of diseases and trial phases, providing a robust foundation for training and evaluation. By leveraging this dataset, *HINT* was able to demonstrate its ability to generalize across different types of trials and diseases (Fu et al. 2022 10-11).

HINT performance was validated through a series of experiments, showing significant improvements over existing models. For phase-level predictions, *HINT* achieved F1 scores of 0.665 for Phase I, 0.620 for Phase II, and 0.847 for Phase III. These results highlight *HINT*

ability to handle the unique challenges of each trial phase, from toxicity testing in Phase I to large-scale efficacy studies in Phase III. *HINT* also demonstrated strong performance across different disease categories. For example, it achieved an F1 score of 0.867 for respiratory diseases, 0.786 for digestive diseases, and 0.585 for oncology. While oncology trials proved to be the most challenging, *HINT* performance in other categories underscores its versatility and robustness (Fu et al. 2022, 14).

Case studies further illustrated *HINT* practical utility. For instance, the model accurately predicted the failure of Entresto's \$200 million Phase III trial, assigning it a low success probability of 0.476. Conversely, it successfully forecasted the outcome of Sitagliptin's diabetes trial, assigning it a high success probability of 0.742 (Fu et al. 2022, 14-15).

Despite its strengths, *HINT* is not without limitations. The model is currently limited to small-molecule drugs, excluding biologics and medical devices. It also struggles to work well in case of rare diseases due to the lack of sufficient training data. Additionally, *HINT* simplifies outcomes into binary success or failure labels, which may not fully capture the nuance of trial results. Ultimately, the complexity of the hierarchical interaction graph can make interpretation significantly more difficult.

Future work aims to address these limitations by expanding *HINT* to other trial types, integrating outcome labels with higher granularity, and improving interpretability through explainable AI techniques. These advancements would further enhance *HINT* utility and applicability. So, the Hierarchical Interaction Network (*HINT*) represents a paradigm shift in clinical trial modelling. By integrating multimodal data, leveraging external knowledge, and employing advanced graph-based techniques, *HINT* sets a new standard for accuracy and scalability in trial outcome prediction. As the pharmaceutical industry continues to embrace AI-driven solutions, *HINT* offers a powerful tool for optimizing drug development and improving global healthcare outcomes.

Even though *HINT* holds a fundamental role in this field, it still grapples with limitations, for instance challenges with handling missing data and fine-tuning predictions for phase-specific regulatory needs. Fu et al.'s. (2022) work has sparked subsequent studies expanding on *HINT* framework by including innovative machine learning techniques, robust data imputation methods, and large language models with the goal to improve predictive robustness and regulatory adaptability in clinical trial setting which will be explored in the following literature review.

Missing values, due to inconsistencies, can undermine the reliability of the success predictions, as they can result in potential biases and reduced modeling robustness. Lo et al. (2019) directly target the data completeness limitation described in Fu et al.'s (2022) work. By enabling accurate predictions even with incomplete data, this study enhances the foundational methods of *HINT*, allowing more robust trial predictions and continuity despite data gaps. Lo et al. (2019) study utilizes advanced machine learning techniques to impute the missing values in their dataset, for instance k-Nearest Neighbor (kNN) and Multiple Imputation by Chained Equations (MICE). Training machine learning models such as Random Forests and Support Vector Machines on the imputed datasets has shown a heightened predictive accuracy for drug approval outcomes. Lo et al. (2019) have allowed for more reliable trial predictions with their work.

While Lo et al. (2019) enhances prediction reliability by improving data completeness, Chen et al. (2024) improve clinical trial outcome predictions by integrating selective classification into the existing *HINT* framework to manage uncertainty. This approach permits the model to refrain from low-confidence predictions while boosting the accuracy of trials with increased prediction confidence. The authors demonstrate that including selective classification (SC) into clinical trial outcome models significantly enhances predictive reliability, which refines *HINT* framework by adding reliability in early, uncertainty-prone trial phases.

Aliper et al.'s (2023) further strengthens outcome predictions by developing a multi-modal AI framework to predict phase transitions based on trial design, target characteristics and omics through transformer-based AI. The study demonstrates the effectiveness of the applied framework, by reaching a 79% accuracy rate. The model not merely refines the general outcome estimations of *HINT* but further offers phase-specific insights.

Following the advancements demonstrated by Aliper et al. (2023) in multi-modal AI, the potential of multi-modal approaches is further elevated by leveraging Large Language Models (LLMs). Zheng et al. (2024) utilize an ensemble of LLMs, their LIFTED framework, to manage multimodal data inputs for clinical trial predictions. By leveraging a Mixture-of-Experts (MoE) architecture, LIFTED adaptively selects the most suitable data source according to trial phase and particular prediction requirements. This enables the model to tailor its focal point, which in return increased predictive accuracy and robustness throughout the entirety of trial stages. This method extends *HINT* core by portraying the layers clinical trial data but also boosts versatility and predictive accuracy by processing diverse data through LLM ensembles.

As Aliper et al. (2023) and Zheng et al. (2014) employ general multi-modal datasets for phase predictions, Qi et al. (2019) aim at a thorough incorporation of target-based drug data with clinical features. With this, they are depicting patient responses and pharmacokinetics (PK) with the goal of predicting individual trial outcomes. This framework demonstrated improved prediction accuracy, which approach fills a gap in *HINT* comprehensive emphasis, by encapsulating detailed patient-level insights.

The usage of machine learning algorithms such as Random Forests and Gradient Boosting cannot only be found in relation to missing value imputation by Lo et al. Kavalci & Hartshorn (2023) also use machine learning algorithms to evaluate trial data and predict early termination based on recruitment rates, compliance, and trial characteristics dynamically. The

study achieved high accuracy in selecting trials with high dropout probability, making room for preemptive refinement. Kavalci & Hartshorn (2023) not only complement the foundational factors of the *HINT* model by tackling with trial completion rates, but they are also able to detect risks prematurely, which consequently improves resource allocation and cost-efficiency in trials. Building on Kavalci & Hartshorn (2023), leveraging machine learning algorithms to predict early trial dropout and manage resource allocation, Ghim & Ahn (2023) further boosts cost efficiency by implementing LLMs into clinical trial operations, as we have seen similarly in Aliper et al.'s (2023) study.

Ghim & Ahm (2023) have undertaken the task of automating clinical trial tasks concerning patient matching by eligibility criteria, protocol creation and explicit consent by taking advantage of AI driven chatbots simplifying complex trial information. The authors demonstrate a significant enhancement in trial efficiency, as patient matching times have reduced, while protocol generation and document processing reap the rewards of LLMs' ability to manage massive quantities of trial data. Additionally, they found that LLM-based chatbots enhanced subjects comprehension by simplifying medical information, emphasizing that LLMs can contribute creating not only operational but also participant-centered improvements in clinical trials. These enhancements were outside of the scope of Fu et al.'s (2022) work but support their objective of enhancing trial success by reducing administrative delays.

In the same vein as Ghim & Ahm (2023), Reinisch et al. (2024) also addressed operational efficiency by predicting which clinical trials will advance from Phase III to final approval. This has been approached by finetuning CTP-LLM to clinical trial documents, reaching high accuracy and reliability. This indicates that CTP-LLM can aid in efficient resource allocation and improved decision-making throughout trial phases. Echoing Aliper et al.'s (2023) work, Reinisch et al. (2024) offer phase-specific insights and thus enriching *HINT* broad success

estimation approach.

Efficient operations alone cannot compensate for the challenges of recruitment, which remains a dominant factor in trial timelines and completion rates. To address this, Bieganek et al. (2022) develop predictive models zeroing in on enrollment success forecasting, allowing for initiative-taking and operationally efficient recruitment planning. Through examining historical clinical trial data, the model predicted enrollment outcomes with enhanced accuracy, while also detecting influential trial characteristics leading to the recruitment's success and enrollment changes at an early stage. This approach not only lessens delays related to under-enrollment, but it also aids to a more seamless trial completion, directly responding to a key factor of trial success highlighted by Fu et al. (2022).

Gayvert et al. (2017) shift the focus to an alternate key determinant in securing trial fulfillment: drug toxicity. This is executed by assessing toxicity risks influencing participant safety and retention. By including drug toxicity prediction leveraging target-based and structural features of compounds, the authors developed the PrOCTOR model, assessing drug properties to measure toxicity risk. The goal of this estimation is to forecast trial success by detecting compounds with manageable toxicity levels. The model displayed enhanced predictive power by differentiating FDA-approved drugs from those that failed due to toxicity with higher accuracy. Furthermore, PrOCTOR was also able to recognize network connectivity and target expression as crucial components in defining a compound's safety profile. Note, that *HINT* does include toxicity as part of the *ADMET* dataset, Gayvert et al. (2017) offer a directed approach to toxicity, delivering a detailed examination regarding patient safety.

As Gayvert et al. examines toxicity, Murali's (2021) dataset includes general bioactivity metrics like drug-target interactions and pharmacodynamic properties. This study moves beyond the safety-focused perspective of toxicity to incorporate molecular-level relations

impacting a drug's efficacy. The analysis highlighted that bioactivity-informed predictions significantly enhance the detection of drugs likely to succeed in clinical trials, specifically in early phases where biological characteristics are critical for predicting efficacy. This study complements *HINT* as it already uses the *ADMET* properties to broadly capture pharmacokinetics and safety. Nonetheless Murali (2021) can offer a deep dive into bioactivity metrics, such as drug-target interactions and pharmacodynamics, providing a more nuanced of a drug's impact.

3. Exploratory Data Analysis

3.1 Data Collection and Description

In this chapter, we are diving into the exploratory data analysis which we conducted on the *ClinicalTrial.gov* dataset to uncover patterns, anomalies and structures, providing insights for our analysis. The dataset is highly relevant for clinical trial research, containing detailed attributes on study designs, phases, outcomes, and interventions.

Key columns like “primary_outcome_measures”, “eligibility_criteria”, and “intervention_model” provide substantial information to support diverse analyses.

The data was sourced from *ClinicalTrial.gov* in the form of a JSON file, which was then transformed into a CSV file for easier analysis and manipulation. The dataset comprises 512,835 rows with 32 columns. However, given that our analysis is merely centered on interventional clinical trials dealing with small drugs and it includes only concluded clinical trials, we disregarded the rest, leaving us with a filtered dataset containing 132,129 rows and 32 columns. From this point onwards, we are referring to clinical trials, whereby we only understand this term to mean interventional studies. Also, the “study_status” only includes clinical trials with four different attributes that represent the trials that are concluded, such as “COMPLETED”, “TERMINATED”, “WITHDRAWN”, and “SUSPENDED”.

3.2 Data Cleaning and Preprocessing

In preparation for a detailed analysis, it is essential to execute preprocessing action to address issues such as missing values and duplicates ensuring robust data quality and integrity.

An initial look of the dataset was performed revealing that there were close to no missing values in any of the columns apart from the enrollment, having 2,304 missing values which will be maintained. Although these missing values would typically be removed in a standard approach to maintain data consistency, we decided to keep them since the column “enrollment” hold a central role in our dataset. A final decision will be made after progressing further, specifically when we are curating the final dataset including this data. Nonetheless, we also found the presence of “NA” values in the “phases” column, which could imply ambiguous classification for some trials and needs to be further investigated for accurate phase specific analysis. However, a closer inspection has uncovered a substantial number of blank entries in several columns, prompting more attention and thereby setting the stage for targeted data cleaning measures. These columns are for instance “collaborators” (69.52% blank), “detailed_description” (39.64% blank), “keywords” (35.89% blank), “maximum_age” (45.40% blank), and “secondary_outcome_measures” (21.51% blank). Given the share of blank spaces (20%) as well as their irrelevance for further analysis, they were discarded. Nevertheless, other columns such as “arms” (11.36% blank), “locations” (9.01% blank), and “primary_outcome_measures” (4.04% blank) also displayed blank spaces. Converting the format of “completion_date” and “start_date” has led to the same amount of missing data in both columns, which suggesting that the transformation has been successful. Upon addressing the missing values, we conducted further checks for data quality issues in regards of negative values and duplicated rows. Our analysis confirmed that the dataset is devoid of negative values and duplicated rows. The following table depicts the number of missing values in our *ClinicalTrial.gov* dataset.

Missing values overview

	Collaborators	Detailed description	Keywords	Maximum Age
# Missing values	87,856	50,095	45,363	57,379
% Missing values	69.52%	39.64%	35.89%	45.40%

Table 1 - The table presents the count of missing values and the corresponding percentage of missing values for each attribute.

Overall, we conclude that the dataset shows strong completeness in crucial areas, due to it being fully populated in key columns such as “ntc_number”, “brief_title”, “study_status”, and “sponsor”, ensuring consistency in essential data attributes. It also demonstrates significant consistency across numerical and categorical data, which is proven by the absence of duplicates and negative values, confirming the integrity of the dataset. Nonetheless, the presence of “NA” values in “phases” indicates potential gaps in classification, which may impact the accuracy. In terms of accuracy we conclude that the dataset structure aligns with expected standards for clinical trials data. Despite this, some descriptive columns with a high number of missing records, such as “collaborators”, are to be carefully analyzed to not bias the model outcome. Consistently populated fields for “start_date” and “completion_date”, enabling chronological analyses of clinical trials suggest the great timeliness of the dataset, ensuring an effective study of trends over time, such as trial durations or historical shifts in trial designs. The dataset’s high usability and relevance is suggested by its substantial information to support diverse analysis through its detailed and structured attributes. After this general analysis of the dataset, we are diving deeper into the contents of the dataset.

3.3 Insights and Patterns

Our effort of preprocessing and cleaning the data has set essential groundwork to explore the underlying patterns and insights, which are vital to understanding the complexities of our dataset.

3.3.1 Study Completion

By examining the historical progression of clinical trials over the years we discovered an increasing trend, starting in the year 2000, which was likely due to a rise in medical research (See Appendix, Graphic 1). This can be measured in multiple ways, although one of them is the rapid growth of scientific articles published in the Directory of Open Access Journal starting in the year 2000 (Laakso & Björk, 2012, 4-6). Additionally, we observed a notable rise in initiation and completion of clinical trials occurring around 2020, which could potentially have been impacted by the global COVID-19 pandemic, causing a massive increase in funding for vaccine research as seen by the investment of Operation Warp Speed with a total of \$18 billion in COVID-19 related research and development (World Health Organization, 2023, 7). Through this longitudinal analysis, we also uncovered that the amount of completed studies is slightly higher than the number of newly started trials, indicating that many recent trials are still in progress. Our investigation into temporal dynamics also gives us insights into completion dates extending into the 2030s, which illustrates how some studies collect data over elongated periods of time. Apart from completion dates being far in the future, we also found starting dates in the future, indicating the long-term project management roadmaps in clinical trials.

Moving beyond the historical insights gathered, we now broaden our discussion to the general landscape of study completion. While our analysis demonstrates a high completion rate of clinical trials across all phases due to most studies being completed, a small but noteworthy amount still failed to fall under this category as these studies have been either terminated or withdrawn. This is likely due to factors such as feasibility issues, safety concerns, or lack of efficacy. However, we can assert that a suspended study is a relatively rare occurrence in clinical trials. With general completion patterns established, it was also worth investigating how completion evolves throughout the year. It became apparent that completed trials showed

a stable trend, while terminated and withdrawn trials were consistently depicted in smaller numbers and suspended trials are notably rare in any month. The month of December is the object of a significant surge in trial completion, which could be explained by year-end finalization. Next to yearly trends, we also narrowed our perspective to the structured phases of a trial. Phase II exhibits the highest amount of completed trials, followed by Phase I, III and IV, respectively. However, we still observed a considerable number of completed trials being unclassified under specific phases ("NA" with 10,235 trials). While phases I and II exhibit the highest amount of suspended and terminated trials, they occur the least in Phase IV. Although withdrawn trials are evenly spread throughout all phases, they occur the most in phases I and II.

3.3.2 Study Duration

After assessing the completion rates, we progress to explore the study duration, which is crucial to understand the efficiency of completion even more. Predominantly the clinical trials in our dataset have relatively short timelines with most of them being not longer than 8 years and the biggest among these trials falling into the duration of maximal 4 years (See Appendix, Graphic 2). Despite these observations, we discovered right-skewed distribution as well as a sizable number of outliers, indicating that a notable number of trials last longer than the average. Additionally, we discovered that suspended trials demonstrated a broader distribution and a higher median duration when compared to other statuses. This is a stark contrast to withdrawn trials, featuring a narrower range and a smaller number of outliers.

3.3.3 Study design

With an understanding of temporal aspects of clinical trials, we now proceed by shifting our focus to explore implications of study designs, to evaluate how different methodologies impact trial completion rates. Randomized studies (67,749) exhibit the highest completion count, followed by non-randomized studies (19,375). While suspensions were rather

uncommon across study types but were slightly more prevalent in non-randomized studies (159), terminations also occur less often in randomized studies (9,227) than non-randomized (4,474). Despite this observation, withdrawals are relatively infrequent but still more notable in randomized (3,869) and non-randomized (1,760) studies. Therefore, our data underscores the effectiveness of randomized trials while indicating potential vulnerabilities to operational challenges of non-randomized studies.

Delving deeper into study designs has also shown us that the parallel model is the most used one while also exhibiting the highest number of completed studies. In this model, randomly assigned patients are given new therapy in the treatment group while control patients are given standard therapy (Cleophas & Vogel, 1998, 113). This is followed by the single group model, which is also widely used despite being impacted by higher termination and withdrawal rates. As already implied by the name, participants are here in only one group and receiving the same intervention. Since there is no comparison group, the outcomes are evaluated by comparing them to baseline measurements (Wang et al., 2024, 1-4). Sophisticated models such as crossover and factorial models are less commonly employed, although generally maintaining reasonable completion rates. While factorial models assess several interventions at the same time by assigning participants to multiple combinations of a treatment (Montgomery et al., 2003, 2), the crossover model assigns both new therapy and standard therapy to participants (Cleophas & Vogel, 1998, 113). Sequential models also see limited usage. The key difference between a sequential and a crossover design is the fact that participants receive one treatment only instead of all treatments in sequential order (Parson et al., 2024, 2-3). Ultimately, the evidence clearly points to a balance between study design complexity and stability with simpler models leading to more frequent utility while also occasionally coming short and consequently leading to early study termination or withdrawal. Given that the choice of the study model can inherently dictate the study's success or failure,

we explored how these models are implemented in randomized and non-randomized studies. Essentially, parallel, crossover and factorial models heavily employ randomized study designs to maintain integrity in their comparative analysis. In contrast, single group models rely on non-randomized settings, due to their ideal utility for exploration or observational studies, characterized by the absence of a control group. Sequential models demonstrate a synthesis of comprehensive blend containing randomized and non-randomized allocations, enhancing their adaptability.

3.3.4 Primary purposes of clinical trials

While the study design lays the groundwork of the study design, it is also pertinent to delve into the primary purposes of clinical trials. Our analysis demonstrated that the dataset is dominated by trials dedicated to treatment (99,833 entries) as primary purpose, suggesting that a notable portion of attention in clinical trials is shifted towards evaluating the efficacy of therapeutic interventions to make progress in patient care. This is followed by prevention studies (7,808 entries), aiming to avoid diseases and conditions. Basic science (6,124 entries) is also prominent and focused on exploring biological processes. As we have defined the primary purpose of clinical studies, this led us to turn to the specific conditions they investigate. Delving into these conditions has shown that “Healthy” is the most frequently studied conditions, which is then followed by "Breast Cancer," "Healthy Volunteers," "HIV Infections, etc. Furthermore, our analysis suggested that clinical trials are heavily focused on grasping drug dynamics (pharmacokinetics), to address health challenges such as HIV or cancer. Most of clinical trials only observed one or two drugs, while the usage of ten or more drugs at once is a rare occurrence, suggesting the desire to focus on simpler regimens. Naturally, “Placebo” was used in most interventions in clinical trials, representing its fundamental role as a control substance. Due to our focus being set on predicting the success or failure of clinical trials regarding the testing of drugs, we discarded all entries titles

“Placebo”. Cyclophosphamide, Carboplatin, and Cisplatin were among other frequently used drugs, which are mostly used in oncology trials, underscoring the significant role of cancer treatment.

3.3.5 Patient Characteristics

After detailing our findings regarding the trial’s purpose, it was imperative to shift our focus to the patient’s characteristics. While majority of the trials includes male and female participants, we found a small amount of gender-specific trials. Upon shedding light on this matter, our data also demonstrates that male-specific trials have the highest success rate, which is followed by female-specific trials and trials involving all genders. This could indicate that gender-specific trials may lead to successful outcomes, due to a more focused study design with clearer treatment effects. Nonetheless, gender-specific trials seem to inadvertently overlook older adults, as female-specific trials include participants between the ages of 21 to 56 years old. In contrast, trials including all genders exhibited a broader age range with ages spanning from 19 to 62 years. Aside from this, there are also several trials dedicated to pediatric conditions, suggested by the exceptionally low minimum ages which are stored in months instead of years. Generally, the minimum age of participants skews towards 18 years.

3.3.6 Enrollment

Patient characteristics directly inform patient enrollment which is why we now turn to exploring this attribute in our dataset. Following our analysis, we found 6,577 records of zero enrollments, which is why we consequently assume that trials with no participants are usually considered as a failure, due to no data being collected. Overall, we also concluded that enrollment values are generally low, suggesting that clinical trials must cope with challenges in recruiting large numbers of participants. After relating enrollment values to clinical trial outcomes, we demonstrated that successful trials (231.77) exhibit a higher average enrollment

in comparison to failed trials (81.74). Consequently, this suggests a positive correlation between higher enrollment numbers and the success probability of trial success. Additionally, successful trials showcase a broader variability in enrollment, as suggested by the higher standard deviation. Our dataset also gives us insights into country-specific enrollment data, which showed that trials were often aimed at addressing widespread health issues such as infectious diseases, which is shown through the high enrollment in developing countries.

3.3.7 Contribution

The list of organizations with the highest number of studies was topped by Pfizer and Novartis, holding 2,442 and 2,436 studies respectively, which are followed by GlaxoSmithKline with 2,161 studies and the National Cancer Institute (NCI) with 1,968 studies. Top sponsors of clinical trials included a diverse mix of universities, pharmaceutical companies, and cancer research institutes, illustrating the broad range of organizations funding clinical trial research. Among these, universities and public health institutions hold a more prominent role in sponsoring. The presence of renowned medical institutions such as the Mayo Clinic and M.D. Anderson Cancer Center underlines the importance of clinical research in progressing patient care and medical innovation. Not only do sponsors hold a special role in advancing clinical trials, but collaborators are also of special consideration. Apart from that, we demonstrated that the National Cancer Institute (NCI) is the one with the highest number of collaborators in clinical trials, indicating their vital role in clinical trials while also suggesting the importance of clinical trials in cancer research. Pfizer and GlaxoSmithKline also frequently collaborate, underscoring their commitment to working with research institutions to advance drug development and clinical research.

4. Benchmark

Our starting point information is the clinical research studies data that we have retrieved from *ClinicalTrials.gov*, which is a publicly accessible website and online database that provides

up-to-date information on clinical research studies conducted across more than 200 countries worldwide.

We have defined and developed an up-to-date standardized benchmark which was essential to build a robust, interpretable and scalable model.

For each clinical trial, we have collected the following five data items: drug molecule information, disease information, trial eligibility criteria, number of enrollment and trial outcome information.

Our benchmark includes molecule information with SMILES of the drugs, the target disease information encoded into ICD-10 codes, the trial eligibility criteria and biomedical knowledge. The encoding into latent embedding vectors of all the data included in the final dataset differs according to their own specifications.-

Standardizing clinical trial data by linking diseases with ICD codes and drug with SMILES (Weininger, 1988) medicine is a common practice. In fact, the approach ensures uniformity in data representation, facilitating the integration of new data from diverse sources and enhancing its utility for machine learning applications.

4.1 Disease data

The disease data are extracted from the collected clinical trial data on *ClinicalTrials.gov* and linked to ICD-10 codes extracted from the World Health Organization Database. The dataset contains disease information including ICD-10 codes and the disease description. ICD-10 is the “international standard for systematic recording, reporting, analysis, interpretation and comparison of mortality and morbidity” data published by the World Health Organization (WHO).

Recently, the ICD-10 has been replaced in January 2022 with the ICD-11 by the World Health Organization (WHO). However, the decision to use ICD-10 instead of the newer ICD-11 in our research was influenced by diverse factors.

Firstly, ICD-10 has a well-established ecosystem of tools, libraries, and datasets, making it the most pragmatic choice in the current time the study was elaborated. Secondly, the ICD-10 codes are still relevant and widely used in past and recent studies in the clinical trial landscape, therefore adapting to past studies guarantees to obtain relevant and interpretable results by the broader scientific and medical communities.

We performed an algorithm which assigns ICD-10 codes to disease names using a heuristic matching approach. It first attempts an exact match, prioritizing precision by checking if the disease name directly corresponds to an ICD code. If no exact match is found, the algorithm performs a partial word match, focusing on significant words with seven or more characters, while limiting matches to simpler disease names to avoid ambiguity.

If these methods fail, the algorithm uses word-overlap matching, comparing the disease input name with candidate names by evaluating shared words and prioritizing those with the greatest overlap and combined word length. This step allows for flexibility in handling variations in terminology. When no suitable match is identified, the algorithm returns no result, ensuring it avoids incorrect associations. This combination of exact, partial, and overlap matching provides a structured approach to map disease names to ICD codes effectively.

4.2 Drug molecule data

The drug molecule data is extracted from the collected clinical trials data on *ClinicalTrials.gov* and linked to the molecule structure, i.e. SMILES string. The SMILES collections with their own respective drug name and drug's description have been collected mainly from the GitHub dataset stored in the *HINT* repository, which retrieved the data from the Drug Bank Database (Fu et al. 2022).

The SMILES, i.e. Simplified Molecular Input Line Entry System, is a notation system that represents the structure of chemical compounds as concise ASCII strings. These SMILES

strings can generate two-dimensional diagrams or three-dimensional models of the molecules (Weininger 2018).

The collected dataset on the Drug Bank Database results as a list of drugs with their respective SMILE molecules. We have performed an algorithm, in the same ways as for the diseases, which maps drug names to their corresponding SMILES representations using a heuristic approach. It first checks for an exact match between the input drug name and a predefined drug-to-SMILES mapping. If an exact match is not found, it searches for partial matches by analyzing individual words within the drug name, focusing on words with seven or more characters that are more likely to yield valid matches. Likewise, the goal is to retrieve their SMILES codes to generalize and standardize the drug data information.

Linking diseases to ICD-10 codes and drugs to SMILES ensures standardization, enabling seamless integration of diverse datasets and supporting the addition of new data in the future. These representations allow machine learning models to incorporate novel information from various sources, enhancing their capacity to test and predict outcomes for emerging drugs or diseases.

4.3 Trial eligibility criteria

The eligibility criteria outline the parameters for selecting participants in a clinical trial, specifying both inclusion and exclusion requirements. They are presented in unstructured natural language, detailing participant characteristics such as age, gender, medical history, current health status, and the target disease or condition. Each clinical trial's eligibility criteria, along with its respective clinical case, are published on *ClinicalTrials.gov*. In this perspective, we separated the inclusion and exclusion criteria and applied language processing to the unstructured text data, preparing it for the encoding phase to generate embedding vector, as later explained in detail.

Eligibility criteria are crucial prerequisites for the success of clinical trials, as they

significantly influence the outcomes. Poorly defined eligibility criteria can result in inadequate participant recruitment, a key factor contributing to the failure of many clinical trials (Su & Cheng 2023).

4.4 Trial outcome information

The TOP benchmark’s target label represents the success or failure of the trial, which indicates whether the trial meets the primary endpoints or not. The labels have been manually curated and are influenced by the feature p-value and why stop present in the XML file on *ClinicalTrial.gov*. Our research instead targets the completion trial prediction which focuses on determining whether the trial would be completed or not. Specifically, the labels are defined based on the clinical trial’s final status. To determine whether the trial would terminate, we excluded all cases where the status indicated an interim step prior to the study’s conclusion. These cases could potentially result in either a positive outcome, such as completion, or a negative outcome, such as termination, withdrawal, or suspension.

4.5 ADMET

The pharmacokinetics (PK) knowledge is leveraged to pretrain embeddings, emphasizing how the body processes drugs after administration. This approach is driven by the critical role that PK factors and disease risks play in determining trial outcomes. Key properties such as Absorption, Distribution, Metabolism, Excretion, and Toxicity (*ADMET*) are central to this process. Integrating these *ADMET* factors into our pretraining ensures a deeper understanding of the interactions between drugs and the human body, enhancing the model’s ability to predict the success or failure of clinical trials.

The different components of the *ADMET* could be described as follows: firstly, the absorption component refers to the process by which a drug is transported from the administration to the site where it will exert its pharmacological effect. Secondly, the distribution component demonstrates the movement of a drug to and from various tissues within the body, as well as

the number of drugs present in these tissues.

Moreover, the metabolism component refers to how drug molecules are broken down by enzymatic systems, thereby influencing the duration and intensity of the drug's action. The excretion component describes the removal of drugs from the body via various excretion routes. Finally, the toxicity component refers to the extent to which a drug may cause harm to the body. An early assessment and prediction of these properties in the drug development workflow optimizes compounds with pharmacokinetics (*ADMET*) characteristics and minimizes toxicity (Vrbanac & Slauter 2017). As a result, it can assist in the prioritization of drug candidates, thus reducing the number of costly trial failures and optimizing resources for the successful development of new drugs.

We retrieved the *ADMET* data from the GitHub dataset stored in the *HINT* repository, which retrieved the data from the Drug Bank Database (Fu et al. 2022).

4.6 Enrollment

The enrolment feature is extracted directly from clinical trial data available on *ClinicalTrials.gov*, distinguishing itself from the TOP benchmark. Given the observed correlation between a high number of enrolments and an increased likelihood of clinical trial success, it has been included alongside other variables due to its potential influence on the model's outcomes. In fact, including a high number of participants poses challenges due to the cost and difficulty of recruiting individuals with the specific characteristics required for clinical studies, which can ultimately influence the likelihood of the clinical trial's success or failure.

Unlike the other variables, which are sequential in nature, the enrolment feature is a numerical, tabular variable. This distinction makes it unique in the dataset, as it does not require preprocessing before being embedded into the model.

4.7 Data preprocessing

As previously illustrated, we filtered the dataset to achieve a complete and clear dataset. In a first step, the number of trials was 512,835 in total. After filtering based on the study's status to determine the study's outcome label, the total number of clinical studies is 325,725.




Figure 3 - Data preprocessing pipeline

Since the model includes the molecule encoding to leverage the potential of the multi-model data source, it must handle only interventional trials involving small molecules, while other trial types such as medical devices and biologics trials are excluded. The total number of interventional studies including only those involving drugs is 132,129. Moreover, we have excluded those rows with missing values in the “phase”, “start_date”, “completion_date” and “drug” features, giving us a total of 107,656 clinical studies. Furthermore, we selected trials with known drug molecule structures and available disease codes, i.e. the trial we were able to match with the integrated datasets of the SMILES and ICD-codes, leaving us with a total of 102,655 clinical trials.


4.8 Data split

To evaluate the predictive model in an effective way, we split the data into three datasets: the training, validation, and test sets. This process has been done chronologically, meaning that we used the column "start_date" as an index so we can have more recent data in the test set. This splitting is following the rule of 70/15/15, consisting of having 70 % of the data in the training set, then 15% in the validation and the remaining 15% in the test set. This is a good approach since most of the data is utilized for the training set while the validation and the test have the same amount. The training set is used for the development of the model, the validation set is used for hyperparameter tuning and to prevent overfitting, while the test set is

reserved for the final evaluation of the model. Since we are going to be focusing on looking at a specific phase, we also split the data according to their phases, leading to us having the data split in “PHASE1”, “PHASE2” and “PHASE3”. We did not include a dataset for “PHASE4”, since its purpose is the monitoring of long-term effects of a treatment, which was already launched into the market according to the National Cancer Institute. Since our work is rather focused on supporting stakeholders to lead with an assertive strategy, to save resources like time and money, “PHASE4” data is out of scope of this work.

Consequently, our approach is to apply a filter for the column “Phases” with the result of only including entry with the value of a specific phase. This caused the data to be mapped to the data with specific phases, which was motivated by the fact that the column still contained a significant amount of data with more information apart from just one phase, such as “PHASE1|PHASE2” (6,155 entries), “PHASE2|PHASE3” (3,016 entries) and “EARLY_PHASE1” (1,834 entries).

The dataset demonstrates a pronounced imbalance in label distribution across all three phases, with a significantly higher proportion of clinical trials labeled as successfully completed compared to those labeled as not completed. Hence, the percentage of successful trials ranges from 78.85% to 85.96%, meaning that over one in five clinical trials is unsuccessful, as illustrated in Graphic 1. This imbalance is particularly evident in the training dataset and has a substantial impact on the model's performance. To mitigate this issue, appropriate evaluation metrics will be employed to ensure a fair and accurate assessment of the model’s predictive capabilities.



Graphic 1 - The table presents the label distribution of clinical trials categorized by their respective phases within the training dataset. The top-left image illustrates the distribution for Phase I, the top-right image corresponds to Phase II, and the bottom.

Additionally, we aimed to address the specific characteristics of each disease, therefore the data was further categorized by disease type. ICD codes had been used to map the clinical conditions with the aim of broadening the disease categories through the CCSR (Clinical Classification Software for Services and Procedures). The new information retrieved allowed us to refine the analysis by clustering similar conditions together, enabling us to create more focused and insightful evaluations according to the attributes of similar diseases and their trial outcomes. The disease areas selected are the nervous system (NVS), neoplasms (NEO), infectious and parasitic diseases (INF), respiratory system (RSP) and digestive system (DIG). These categories were chosen due to their significant impact on public health and because of their high frequency in the trial's dataset. This data has been split following the same process

implemented to split the total dataset.

5. *HINTBasic* model

The *HINTBasic* model is a very sophisticated architecture design which was created to predict the clinical trial outcomes by integrating several data sources like diseases, drugs and trials protocols. This model is going to process and encode this information before integrating it into the final framework. Each feature is going to contribute in a unique way to capture the very complex relationship between the variables that define if a clinical trial will reach the end of their phases.

To include more medical information, we included a dataset of ICD codes. This means that we mapped these codes to their corresponding Clinical Classifications Software Refine (CCSR) categories. These codes are hierarchical, which means that they represent a relationship between diseases in a broader category which were then split into more specific categories: For example, the more general code "G90521" can also be split into variations such as "G9052", "G905" and "G90". This marked a crucial step in our modelling process, since it caused a reduction in the diversity of our data, which consequently enabled a division of clinical trials based on diseases and not the high quantity of conditions. For the trials with multiple ICD codes, we decided to create a one-hot vector to create a single composite vector that captured all related medical conditions. This approach allowed the model to take advantage of the hierarchical structure of medical conditions, enabling it to link trial outcomes with participants health conditions more efficiently. In this part we used a graph-based attention model, called GRAM, which is going to encode these hierarchical relationships. This did not only create embeddings for each ICD code, but it also captured the semantic similarities between related diseases. As a result, the model's performance is enhanced due to the added medical information. This is very important because it helped the model understand how certain diseases might influence the clinical trials outcomes, such as having a specific

rare disease will decrease the probability of success, since it is very challenging to find eligible participants.

The protocol encoder addressed the complex text description that is in the eligibility criteria. This column includes information associated with inclusion and exclusion criteria, which gives insights into the criteria that patients need to have or not to be an eligible participant of the clinical trial. For example, trials testing for drugs to cure cancer only consider patients which not only have a specific cancer but are also in a specific phase of the disease. It is crucial to test the effectiveness of the drug in the treatment, since trial protocols are essential to determine the scope and the focus of the trial's proposal. To get all the information about the inclusion and exclusion, the data was split into two criteria, the inclusion and the exclusion criteria. After that, we employed Clinical Bert to generate the encoders at the sentence level (Guangyu Wang 2023). The protocol embedding used the neural embedding techniques to analyze the descriptions, applying layers of processing that contained highway networks, to extract and refine the semantic meaning of the text. By converting text data into high-dimensional embeddings, the protocol encoder ensures that the small details of trial design are accurately reflected in the model, providing valuable insights into how trial structure impacts its success or failure.

The molecular encoder is designed to handle data that is related to the chemical compounds, such as drugs that are being tested in the clinical trial. This encoder can capture the structural and chemical properties of the molecules and then transform that information into embeddings so it can be utilized in the model. This transformation is essential since the molecular properties of a drug significantly impact its safety, efficiency and the potential interactions within the trial. One illustrative example of the potential risks associated with clinical trials revolves around the drug "Trastuzumab Deruxtecan". Despite the occurrence of several adverse effects, including pneumonitis, which has resulted in fatalities, the drug was

ultimately approved following a comprehensive evaluation of its safety profile (Modi et al. 2022). This encoder uses the Message Passing Neural Network denominated MPNN, which is a type of graph-neural network designed to process the molecular structure. In the MPNN, molecules are represented as graphs, which depict atoms as nodes while the bonds are illustrated as edges. As a result, it can reflect the inherent molecular chemical properties and structures. This is crucial for capturing the relationships and the interactions between different atoms. The MPNN is divided into two phases, with the first one being the passing message and the second one being the readout. This means that during the propagates the information across the graph by exchanging messages between nodes (atoms) and updating their representation based on the properties of their neighbors and the edges that are connected to them. This repetitive process enables the model to capture the local features, that are the type and properties of a specific atom, and the global features such as the connectivity of the entire model. In the second phase the molecular encoder is aggregating the updated node representations into a single vector to represent the entire model. This aggregation was done to ensure that the embedding captures all the chemical and structural of the drug. Embedding is of high importance as it enables the model to contain a rich representation of molecular properties, influencing the drugs behavior in the human body. This is due to its ability to provide information on how a drug might perform in a clinical trial, particularly regarding safety and efficacy.

The data loader plays an essential role in efficient preprocessing and batching of datasets for clinical trials analysis and *ADMET* prediction tasks. The dataset structure for clinical trials manages multi-modal clinical data. It contains features such as NCT ID, labels indicating the trial outcome and SMILES to represent the molecular compounds. Additionally, it also contains ICD-10 codes for disease identification, inclusion and exclusion criteria, and an additional feature, the enrollment number in its second version. The framework for *ADMET*

focuses on predicting binary outcomes for *ADMET* properties, with SMILES strings as input features, and their corresponding label are binary (1 for positive and 0 for negative).

In addition, collation methods were used for batching and model training. In the context of *ADMET* prediction, the SMILES string and their corresponding labels are grouped together in batches, while labels themselves are converted into machine learning compatible tensors. This includes encoding ICD codes, protocols, and molecular structures. For instance, SMILES strings are transformed into structured lists and transforming text-based criteria into numerical feature vectors. All preprocessing steps are designed to ensure that the data is clean, organized, and suitable for our downstream tasks.

The data loading pipeline represents a crucial component of the machine learning workflow. It is responsible for preparing clinical trial data from CSV files and structuring it into training and evaluation batches. For *ADMET* prediction, it processes training and validation data for each property, creating task-specific datasets. This modular approach optimizes resource utilization of resources through efficient batching and merging of the data, while transforming data into structured, model-ready input, providing a robust foundation for machine learning models.

Combining *ADMET* with clinical trial data enables researchers to design more targeted and informed clinical trials, increasing the probability of successful drug development. Consequently, development is streamlined with predictive models that enhance clinical decision-making are supported. This also causes improved drug efficacy and safety and reduces costly clinical trials failures.

In the Graph Neural Network, the layers are designed to learn the interactions between different encoded features. Each node in the GNN represents an encoded feature like a disease, protocol information or a molecular property, while the edges represent the relationships between the features. Through iterative aggregation the GNN layers improve the

node representation by adding more information from the neighbors that they are connected to. This process allows the model to build a comprehensive view of the relations between several trials' factors enabling the model to make better informed predictions.

Highway layers in the module are used as a mechanism to allow a smoother gradient flow during the training, which is crucial for preserving information across deep neural networks. These layers operate by applying an adaptive computation approach, where each transformation is balancing the nonlinear and the linear way. These specific layers calculate the weight combination of a non-linear transformation and a similar linear transformation. These weights are determined by a gating function that uses sigmoid activation to control the flow of information, which ensures that the essential features from earlier layers are retained while allowing the network to still learn complex patterns. Adding these layers is highly beneficial for the model since it will prevent the vanishing gradient. This is since holding that specific gradient creates more complex training in deep networks, while also enabling the model to refine representations effectively.

The model also uses the graph attention network that is denominated as GAT, to enhance the learning process by applying the attention mechanism to the graph. Unlike the other standard GNNs, GAT assigns the weights dynamically to edges based on their relevance, which allows the model to focus on the most important connections in the graph. Having this flexibility is very valuable since it can capture subtle interactions between clinical trials elements.

Joining the Highway layers and the GNN layers creates a cohesive framework that works inside the *HINTBasic* model. The highway refining the features during the early stage of processing can ensure efficient information flows and prevent the loss of critical data.

The model's primary focus is on its interactions between the clinical trials protocols, molecular data and the disease information using the advanced neural network components to produce meaningful predictions. After, the model initiated the encoding for each data point.

This process raised for the protocol encoder regarding the text data of the eligibility criteria. The molecular encoder processed the chemical structure of the molecules, while the ICD code encoder operates regarding the hierarchical disease classifications. These encoders transformed the raw data into high dimensional embeddings that captured the essence of the information. Once the embeddings are generated, these are combined into an interaction representation, encapsulating the relationships between trial elements.

A critical step involves computing a “disease risk” embedding, which represents the likelihood of complications related to the target condition. This risk embedding is then integrated with the interaction representation in the “augmented interaction” node, producing a richer and more informative feature set.

Pharmacokinetics (PK) node leverages the *ADMET* framework, each property is modelled independently, and the resulting embeddings are combined into a unified PK representation, capturing the drug’s pharmacological profile.

The enriched interaction embedding from the augmented interaction step is combined with the PK embedding to form a trial-level representation. This trial embedding is processed through a final node, which generates the predicted label, showcasing the probabilities of the clinical trials reaching the end of their phases, indicating their success or failure. Since the model employs a multi-task learning framework this allowed us to simultaneously optimize the model during the training, while also taking the used features into consideration.

6. *HINTPlus* Model

Enrollment is one of the mains factors leading to failure of clinical trials. To include this significantly impactful factor, we included the "enrollment" feature in our model, as this adds more crucial information about the success and failure of clinical trials (Kim et al. 2023)

It was necessary to process the “enrollment” feature in the data loader to create a complete clinical trial version with one additional variable compared to the original dataset. In the

HINTPlus model the "enrollment" feature was treated as node on the GNN layers just like the other features in the previous *HINTBasic*. This allows the model to interact with all the features (ICD code, protocol and the molecular encoders) capturing the new relationships created by adding the enrollment.

In the *HINTPlus* model, the “enrollment” feature was added and shared across multiple tasks, enabling the model to learn how enrollment size influences a trial’s outcome. This integration ensures that the model incorporated an additional important feature, which contributes meaningfully to predictions and, in turn, enhanced the model’s ability to predict success or failure.

Furthermore, this model has been adapted to create specific versions according to the dataset featuring some selected diseases. The primary difference between the *HINTBasic* model and the specialized *HINTPlus* is to be found in the dataset specificity. This adjustment has been made to assess whether the specialized *HINTPlus* shows an improved performance than the generalized one. The following pipeline depicts the working processes of the developed models:




Figure 4 - Pipeline of *HINTBasic* and *HINTPlus*

7. Results

Classical machine learning methods, such as Random Forest and LogisticRegression, and traditional deep learning models like Multi-Layer Perceptron's, often rely on limited features and fail to capture complex interactions among trial components. In contrast, the *HINTBasic* model leverages multi-modal data sources, integrates deep interaction mechanisms through weighted nodes, and incorporates pretraining with the *ADMET* dataset. This approach allows *HINTBasic* to dynamically model complex relationships, prioritize critical features, and utilize domain-specific drug knowledge, resulting in more accurate and reliable clinical trial outcome predictions.

We tested three different models: we focused on traditional predictive models, such as the *Logistic Regression*, the *XGBoost* (Chen & Guestrin 2016) and modern neural network models, such as the *Multi-Layer Perceptron*. The models have been tested using the PyTrial, which is a “Python package providing benchmarks and open-source implementation of a series of ML algorithms for clinical trial designs and operations” (pytrial.readthedocs.io) including trial outcome predictions. It includes an evaluation of clinical trial outcome predictions using the TOP benchmark (Fu et al. 2022) involving different machine learning prediction models and the *HINT* itself (Wang et al. 2023).

The evaluation metrics to compare the diverse models are ROC - AUC, F1 and PR-AUC for evaluating the binary classification, i.e. the trial outcome labels. Those three metrics are widely used to evaluate the performance of predictive models, for classification tasks as in our case. Additionally, given the unbalanced nature of our dataset, F1 and PR-AUC are particularly well-suited as they are designed to handle imbalanced data by focusing on precision, recall, and the performance of the minority class.

The metrics we used are detailed, as follows: firstly, the Receiver Operating Characteristic – Area Under the Curve (ROC - AUC) measures a model's ability to distinguish between positive and negative classes. The ROC curve plots the true positive rate (sensitivity) against

the false positive rate (1-specificity) at various thresholds. The AUC (Area Under the Curve) summarizes this curve into a single value, where 1 indicates perfect discrimination and 0.5 represents no better than random guessing. ROC-AUC is popular for its threshold independence, making it suitable for evaluating models when the relative costs of false positives and false negatives are unclear (Fawcett 2006; Huang & Ling 2005). Secondly, it balances precision and recall, providing a single measure of a model's accuracy that accounts for both false positives and false negatives. Then, the F1 score is particularly valuable in situations with imbalanced datasets (Van Rijsbergen 1979; Sokolova & Lapalme 2009). Lastly, the Precision-Recall Area Under the Curve (PR-AUC) evaluates the relationship between precision and recall across thresholds, focusing specifically on the performance of the positive class. Unlike ROC-AUC, PR-AUC ignores the true negative rate and is particularly suited to imbalanced datasets, where the positive class is underrepresented. PR-AUC provides a clearer picture of a model's ability to capture positive cases, especially when false positives and false negatives carry significant costs (Davis & Goadrich 2006; Saito & Rehmsmeier 2015)

7.1 Evaluation of Our Models

To assess the efficacy of the *HINTBasic* models in diverse scenarios, we conducted experiments on both single phases (Phase I, Phase II, and Phase III).

The two models adhere to the same structured methodology, involving data preparation, model initialization, training, and evaluation. About the single-task *HINTBasic* model, each phase was processed independently, with datasets partitioned into training, validation, and test sets. Prior to training, the model was pre-trained on relevant data, thus enhancing its capacity to capture characteristics specific to each feature.

7.2 Comparison between *HINTBasic* and *HINTPlus*

Results comparison between <i>HINTBasic</i> and <i>HINTPlus</i>						
	Phase I		Phase II		Phase III	
	<i>HINTBasic</i>	<i>HINTPlus</i>	<i>HINTBasic</i>	<i>HINTPlus</i>	<i>HINTBasic</i>	<i>HINTPlus</i>
PR-AUC	0.717 ± 0.007	0.878 ± 0.005	0.645 ± 0.007	0.783 ± 0.007	0.745 ± 0.007	0.822 ± 0.007
F1	0.835 ± 0.005	0.931 ± 0.003	0.784 ± 0.005	0.865 ± 0.005	0.854 ± 0.005	0.899 ± 0.005
ROC-AUC	0.724 ± 0.008	0.878 ± 0.007	0.547 ± 0.006	0.846 ± 0.007	0.535 ± 0.010	0.807 ± 0.010

Table 2 - Results of *HINTBasic* and *HINTPlus*

The inclusion of the enrollment feature consistently improved the performance of the *HINT* model across all three phases. An examination of the PR-AUC and ROC-AUC metrics, which assess the model's capability to differentiate between successful and unsuccessful trials, revealed enhancements across over each phase when comparing the two models. Furthermore, the F1 score, which balances precision and recall, exhibited gains across phases, indicating that the integration of the enrollment feature resulted in enhanced prediction quality.

The model demonstrated optimal performance in Phase I and Phase III, with or without the enrollment feature. Although, it outperformed across all three phases, achieving robust results with the inclusion of the enrollment feature. Precisely, it achieved F1 scores of 0.931, 0.865, and 0.899 in Phases I, II, and III, respectively. The inclusion of the enrollment feature consistently enhanced performance across all metrics, demonstrating its critical role in improving predictive accuracy. Overall, the enrollment feature proved to be a valuable addition, providing meaningful information that enhanced the model's ability to generate more accurate predictions.

7.3 Evaluation of the *HINTBasic* with baseline models

We trained the models using our up-to-date benchmark dataset, employing the pre-split dataset created earlier for consistent evaluation.

The tables provided represent the performance of four different models (LogisticRegression, XGBoost, MLP, and *HINTBasic*) across three clinical trial phases (Phase I, Phase II, and

Phase III), evaluated using metrics such as ROC-AUC, F1, and PR-AUC.

7.3.1 Phase I

The table represents the results of the mentioned metrics according to the five different models ingesting the clinical trials in Phase I. *HINTBasic* outperforms all other models across all metrics, achieving the highest values for ROC-AUC (0.717), F1 (0.835), and PR-AUC (0.724).

Phase I Baseline Results				
	Models			
	LogReg	XGBoost	MLP	HINTBasic
ROC - AUC	0.608 $\pm$ 0.000	0.615 $\pm$ 0.000	0.599 $\pm$ 0.000	0.717 $\pm$ 0.007
F1	0.764 $\pm$ 0.291	0.763 $\pm$ 0.269	0.765 $\pm$ 0.260	0.835 $\pm$ 0.005
PR - AUC	0.7118 $\pm$ 0.008	0.7121 $\pm$ 0.008	0.7105 $\pm$ 0.007	0.724 $\pm$ 0.008

Table 3 - Table of results of the clinical trials prediction on Phase

7.3.2 Phase II

The table represents the outcome results of the mentioned metrics according to the five different models ingesting the clinical trials in Phase II. Likewise *Phase I*, *HINTBasic* remains the best-performing model, showing significant improvements in all metrics. It achieves 0.645 in ROC-AUC, 0.784 in F1, and 0.547 in PR-AUC. The performance gap between *HINTBasic* and the other models is especially pronounced for ROC - AUC, further emphasizing its ability to handle unbalanced datasets effectively.

Phase II Baseline Results				
	Models			
	LogReg	XGBoost	MLP	HINTBasic
ROC - AUC	0.523 $\pm$ 0.000	0.523 $\pm$ 0.001	0.523 $\pm$ 0.002	0.645 $\pm$ 0.007
F1	0.748 $\pm$ 0.246	0.748 $\pm$ 0.231	0.748 $\pm$ 0.273	0.784 $\pm$ 0.005
PR - AUC	0.589 $\pm$ 0.000	0.596 $\pm$ 0.004	0.599 $\pm$ 0.021	0.547 $\pm$ 0.006

Table 4 - Table of results of the clinical trials prediction on Phase II

7.3.3 Phase III

The table represents the results of the mentioned metrics according to the five different

models ingesting the clinical trials in *Phase III*. In the final phase, *HINTBasic* continues to slightly outperform the other models with ROC-AUC (0.745), F1 (0.854), and PR-AUC (0.535). The consistent improvement in these metrics across phases indicates that *HINTBasic* maintains its reliability as the dataset grows larger and more complex.

Phase III Baseline Results				
	<i>Models</i>			
	LogReg	XGBoost	MLP	HINTBasic
ROC - AUC	0.523 $\pm$ 0.000	0.519 $\pm$ 0.000	0.528 $\pm$ 0.000	0.745 $\pm$ 0.007
F1	0.883 $\pm$ 0.274	0.832 $\pm$ 0.254	0.834 $\pm$ 0.251	0.854 $\pm$ 0.005
PR - AUC	0.416 $\pm$ 0.003	0.414 $\pm$ 0.003	0.515 $\pm$ 0.003	0.535 $\pm$ 0.010

Table 5 - Table of results of the clinical trials prediction on Phase III

7.4 Disease groups evaluation

Similarly to past studies (Fu et al, 2022), we have included an evaluation of the model's performance on different datasets separated per groups of diseases. Given that the best performing model is *HINTPlus*, we present the results in Table 6.

Disease groups evaluation with HINTPlus			
Diseases Groups	<i>HINTPlus</i>		
	ROC - AUC	F1	PR - AUC
DIG - Phase_I	0.8571 $\pm$ 0.0196	0.8405 $\pm$ 0.0133	0.7251 $\pm$ 0.0195
DIG - Phase_II	0.7777 $\pm$ 0.0293	0.8546 $\pm$ 0.0116	0.7464 $\pm$ 0.0176
DIG - Phase_III	0.7747 $\pm$ 0.0321	0.8382 $\pm$ 0.0174	0.7219 $\pm$ 0.0261
NVS - Phase_I	0.7760 $\pm$ 0.0374	0.8297 $\pm$ 0.0262	0.7388 $\pm$ 0.0374
NVS - Phase_II	0.8309 $\pm$ 0.0235	0.8668 $\pm$ 0.0182	0.7811 $\pm$ 0.0273
NVS - Phase_III	0.8469 $\pm$ 0.0362	0.8511 $\pm$ 0.0181	0.7413 $\pm$ 0.0275
NEO - Phase_I	0.7754 $\pm$ 0.0360	0.8294 $\pm$ 0.0213	0.7382 $\pm$ 0.0314
NEO - Phase_II	0.8349 $\pm$ 0.0221	0.8687 $\pm$ 0.0122	0.7838 $\pm$ 0.0162
NEO - Phase_III	0.8338 $\pm$ 0.0430	0.8371 $\pm$ 0.0232	0.7205 $\pm$ 0.0344
INF - Phase_I	0.7486 $\pm$ 0.0338	0.8927 $\pm$ 0.0087	0.8190 $\pm$ 0.0150
INF - Phase_II	0.7783 $\pm$ 0.0181	0.8854 $\pm$ 0.0074	0.8092 $\pm$ 0.0118
INF - Phase_III	0.7730 $\pm$ 0.0231	0.9109 $\pm$ 0.0068	0.8425 $\pm$ 0.0126
RSP - Phase_I	0.8653 $\pm$ 0.0739	0.8504 $\pm$ 0.0483	0.7428 $\pm$ 0.0712
RSP - Phase_II	0.8131 $\pm$ 0.0525	0.7137 $\pm$ 0.0419	0.5565 $\pm$ 0.0500
RSP - Phase_III	0.8326 $\pm$ 0.0854	0.8408 $\pm$ 0.0374	0.7271 $\pm$ 0.0549

Table 6 - Table of results of the clinical trials prediction on HINTPlus per diseases

If the metrics achieved by the *HINTPlus* across specific diseases are compared with the ones

accomplished by the general model, it can be denoted the superior performance of the disease specific one. For instance, if the infectious diseases (INF) are considered, the model shows an improved of F1 score for Phase III and an enhanced PR-AUC across all phases, suggesting that model may return less false positives results (a false positive occur when a trial is mistakenly predictive as a success, while it is a failure).

Similarly, in the cases of nervous system diseases (NVS in Phase II and III), neoplasms (NEO in Phase II and III), respiratory diseases (RSP in Phase I and III) and digestive system diseases (DIG in Phase I and III), a better ROC-AUC is observed.

These results demonstrate that the diseases specific model maintains robust accuracy in the predictions of trials outcomes across the diverse phases despite the degree of complexity of the relationship established.

This performance early confirms the reliability of *HINTPlus* as an analytical tool, making it suitable even for highly specific, complex scenarios like clinical trials focusing on individual diseases.

7.5 Results overview

Our models have been compared to the *HINT* model from the original paper, as well as to other baseline models, to better assess the quality of our approach.

HINTBasic consistently outperforms all other baseline models across all phases and metrics, demonstrating its superiority in clinical trial outcome prediction. Its ability to handle unbalanced datasets is evident in the significantly higher F1 and PR-AUC scores. Traditional models like LogisticRegression and XGBoost lag significantly in performance, as they do not include pretraining with the *ADMET* dataset, which limits their ability to incorporate domain-specific drug-related properties into their predictions.

Furthermore, the *HINTBasic* achieves better results compared to the original *HINT*, primarily due to the upgrades operated on the encoders. Specifically, the protocol encoder has been

enhanced by utilizing the *ClinicalBERT* model (Wang et al. 2023), which has enabled more accurate and effective encoding of the clinical trial protocols. Indeed, it demonstrates the significant advantage of improving encoders to boost performance in prediction models that rely on textual data.

The *HINTPlus* model, with the addition of the enrollment feature, demonstrates relevant improvements over the *HINTBasic* results. It is to be mentioned that the trend observed in *HINTBasic*, where Phase III and Phase I performed better than Phase II, is also evident in the models we developed. This consistency highlights potential intrinsic differences in trial characteristics across phases that influence the predictive framework performance. Notably, the *HINTBasic* model can overcome most of these challenges, turning it into the most effective approach for predicting outcomes across each phase.

Furthermore, it is noteworthy that the *HINT* paper used the *p_value* and *why_stop*, as well as IQVIA, to categorize outcomes as either successful or unsuccessful, whereas our model relies on a different labeling method developed by us. This distinction in criteria may impact results and partially explain performance variability between our models and those reported in the *HINT* model.

To conclude, these findings illustrate that *HINTPlus* model is our best model, as it clearly showcases the value of implementing tailored approaches and feature enhancements in the field of clinical trial prediction. This is further emphasized by the model behavior on preselected diseases, as argued above, as the *HINTPlus* demonstrated meaningful efficacy by adapting to the distinctive characteristics of each disease, resulting in enhanced predictive accuracy across diverse trial phases and complexities.

8. Discussion

Our work's foundation has been data. Its quality and comprehensiveness have highly influenced our outcomes. However, our data has also limited our work. To provide a balanced

interpretation of our findings, we are going to discuss the limitations associated with the data and models used. Additionally, we are going to provide a cost-benefit analysis to measure the model's effect. After that we will dive deeper into the implications of the model's implementations for involved stakeholders.

8.1 Data-related limitations

In terms of data-related limitations, dependency on high-quality and complete data is a central limitation. Due to our model having a multi-modal data architecture, which is leveraging ICD codes, clinical protocols and molecular data, it is highly sensitive to missing values. As a result, this could lead to errors and thus to a significant decrease in predictive accuracy. The fact that our work is based on a neural network model makes it challenging to interpret and validate the reasons for the results. This might also have implications when it comes to changing management, as stakeholders could potentially not have enough trust in the model's accuracy, due to its complex nature and their limited knowledge (Chopra et al. 2024)

Additionally, our work is focused on interventional clinical trials testing drugs. Since we are not limiting prediction on a specific disease, it might fall short due to not considering specific factors influenced by a disease. As a result, our model might lack granularity to predict the clinical trial's success.

This limitation is closely connected to another one we are encountering, which is data scarcity. Although we are working with multi-modal data architecture, it is still possible that our model still does not meet the necessary amount of data to make an accurate prediction. This is even more prominent when considering that our dataset is imbalanced, having more success cases (83,360) than failure cases (19,295). This could happen when researchers want to predict a clinical trial outcome which concerns a rare disease. Due to insufficient data about rare diseases, our model might have to face difficulties since it is not able to learn from robust patterns (Nestor et al. 2018).

8.2 Cost/Benefit analysis

After addressing data-related limitations, it is now equally important to evaluate the economic feasibility of our proposed model. This is done by assessing both costs and benefits to see if stakeholders would benefit from implementation. This calculation's purpose is to provide critical insights into the model's value.

8.2.1 Benefits

The predicted results demonstrated by our *HINTPlus* model can be leveraged to improve the overall efficiency of the trial process. Indeed, one of the main benefits consists in shortening the trial timeline.

Given the integration of this model into a trial streamline, the reduction of the number of failed trials can be assumed to range from 10 % to 15 %. This estimation is achieved through a calculation, broken down as follows.

Given the results in terms of F1 score, observed in Phase III for both the *HINTBasic* and *HINTPlus*, it can be stated that the result achieved by the latter shows better differentiation between successful and failed trials.

Thus, the improvement pursued can be quantified through the following:

% Improvement *HINTBasic* vs. *HINTPlus*

	Phase I		Phase II		Phase III	
	HINTBasic	HINTPlus	HINTBasic	HINTPlus	HINTBasic	HINTPlus
F1	0.835 ± 0.005	0.931 ± 0.003	0.784 ± 0.005	0.865 ± 0.005	0.854 ± 0.005	0.899 ± 0.003
% Improvement	11,50%		10,33%		5,27%	

Table 7 - Improvements in percentage of *HINTBasic* compared with *HINTPlus*

HINTPlus demonstrates an ability of distinguishing successful from unsuccessful higher than 11,50%, 10,33% and 5,23%, respectively to their phases, in comparison to the *HINTBasic*. The comparison has been based on the F1-score metric, as it is the most suitable method to assess the general performance in the case of highly imbalanced datasets, like the one here

employed.

Since real-world factors such as operational complexity, uncertainty and regulatory implications need to be considered and given that multiple and reliable sources estimate AI improvements on clinical trials to be in the range from 10 % to 30 % (McKinsey 2023), we will assume a net 10 % to 15 % reduction of failed trials. Taking this into consideration, the overall duration may be reduced by 15 % to 30 % (McKinsey 2023).

Specifically, patient recruitment and enrollment stages would be more targeted, and the process would be consequentially shortened. Predictive accuracy may also improve the decision-making process, leading to faster Go/No-Go decisions, which will additionally decrease timelines. If looking at Phase-specific statistics, earlier Phases trials (Phase I and Phase II) could feature time-savings spanning between 6 and 12 months; similarly, Phase III trials could benefit from a reduction approximately equal to the 20-25 % of their average duration (2-4 years).

Hence, if a total development timeline of 15 years is assumed, the model integration could lead to a time saving estimation spanning from 1.5 to 4.5 years.

This will eventually ensure improved efficiency, consequently reducing the operational burden on the workflow. Furthermore, it might be helpful to tailor specific strategies to avoid early trial termination and to enhance patient recruitment.

8.2.2 Lower costs implications

In terms of cost, given the benefits analyzed above, potential savings may be computed by making some assumptions. Assuming that 10 %-15 % of failed trials could be early detected and avoided across all phases, which could potentially save \$ 260 to \$ 390 million (10 % to 15 %) per each drug development process, if the average \$ 2.6 billion expenditure is taken into account. To breakdown this estimate, we can make the following calculations:

Patient recruitment accounts from 30 % to 40 % of the overall trial budget expenditure (Getz

et al. 2016), therefore if the model is applied, the cost can be reduced by 15 % at least (e.g. if the cost is equal to \$ 15 million, \$ 2.25 million could be saved). Phase-specific expenditure span from \$ 4 to \$ 52 million (DiMasi et al. 2016), depending on the Phase stage; if an average cost of \$20 million is considered, savings could be estimated in \$3 million (15 %) per phase. Other operational expenses, such as data handling, logistical and outsourcing support, and staff salaries account for almost \$1 million per month (Sertkaya et al. 2016), equal to 12 million per year; thus, a timeline reduction of 2 years can reduce the operational charges for up to \$24 million.

Despite the positive metrics achieved, our *HINTPlus* model can make errors and provide unreliable answers. Specifically, the error rate is estimated to be around 0.101 (1-F1 score = 1-0.899).

8.3 Business implications based on stakeholders

Building on this calculation, it is also crucial to consider all the potentially involved stakeholders in the usage of our model. Our goal is to examine implications, while also providing actionable insights to support informed decision-making in a real-world setting. The following stakeholders are the focal point of this discussion: Researchers, Pharmaceutical Companies, Regulators, Healthcare Providers and Patients.

Before delving into the potential impact of our work regarding stakeholders, the implication represented by adoption costs need to be addressed; for instance, embracing a machine learning tool into a clinical trial workflow brings numerous limitations, ranging from technology costs to legal allegations. Therefore, adopters might face substantial expenses in terms of resources such as technology acquisition, infrastructure set up, trained personnel recruitment, operational recurring costs and possible outsourcing costs (McKinsey 2023).

Hence, the initial investment to adopt a ML-based model is undeniably resource-intensive, even though it may be argued that the payback period is short, as benefits can be observed in

the immediate period following the integration.

Secondly, another implication that has to be highlighted lies in the fact that collaboration and communication across all stakeholders are of significant importance when it comes to successfully implementing machine learning methods into clinical research. An open culture for benefits and drawbacks has to be embraced, while best practice methods need to be shared in order to successfully apply these methods (Weissler et al. 2021, 11-12).

The integration of our model could impact researchers and health care providers work significantly, as it could potentially enhance the success and efficiency of clinical research. The resource optimization that comes with the implementation of our model would also positively impact their work, as the prediction of success or failure would maybe prevent them from starting a clinical trial which is not going to be successful according to our prediction.

However, researchers should also be alerted when it comes to the data used in training the model. To prevent overfitting from happening, which in return causes a poor performance of the model, it is of high importance that researchers validate their models in different settings. This practice ensures that the model generalizes well and is not merely capturing noise in the training data, while not fully considering underlying patterns of the data (Kappen et al. 2018, 4).

The accelerated timelines of drug development do not only benefit researchers but also pharmaceutical companies, as they are able to launch new drugs to the market faster, improving their position in the market as well as their competitiveness and thus potentially increasing their ROI. This is especially vital, since it has been found that AI can significantly reduce drug development. The optimized resource allocation, which comes with the use of AI, can lead to efficient trial design, which impacts savings made in the process of drug development (Weissler et al. 2024, 4-5). The economic impact of drugs being launched into

the market on pharmaceutical companies can be measured in a company's stock prices. The company "BridgeBio Pharma" has seen a significant surge in stock after the FDA has approved their drug based on their clinical trial results, treating a rare heart condition, which might influence market dominance. The company is expected to reach global sales of \$ 2.5 billion by 2035 (Investor's Business Daily 2024).

This stands in stark contrast to the company "Cassave Science", which experienced a failed clinical trial, focused on testing an Alzheimer's drug. The announcement was followed by a stock drop of 84 %, indicating that the company's value was highly influenced by the outcome of this clinical trial. Although pharmaceutical companies can highly benefit from the integration of AI in clinical trials, this also comes with limitations. These include data privacy concerns, regulatory and ethical constraints (GlobeNewswire 2024).

This leads us to examine implications and limitations of regulators. The usage of machine learning in clinical trials is highly relying on evaluations or FDA representatives. These suggest that ML practices in clinical trials are seen as high-risk use cases. The reason for this classification is the potential of errors and biases which are influencing the algorithm. While pharmaceutical companies and researchers are limited by regulators such as the FDA, the FDA is limited by the lack of an existing guideline. However, they are open to collaborating with sponsors and stakeholders on a case basis to support implementations (Weissler et al. 2021, 10-11). Implications for patients would firstly be a minimized treatment burden, as researchers already know beforehand that a trial is not going to succeed, which then prevents unnecessary treatments or therapies from happening. This has the potential to save patients from being impacted by the trial physically and/or emotionally. Additionally, prediction of clinical trial outcomes is linked to early identification of risk factors, leading to adjustments of clinical trials which in return could improve overall health outcomes. Since ML implementations also have a positive effect biomedical evidence, they could potentially save

humans and reduce their suffering (Kappen et al. 2024).

Future research can focus on addressing these challenges by developing solutions using other frameworks, which could be tailored to regulatory and ethical requirements, while also having higher interpretabilities for all stakeholders. While generalizability and scalability can be improved by training the models on more diverse datasets, which include an even wider range of data, it should also be considered that disease-specific models could be able to make an enormous contribution to the prediction of clinical trial outcomes.



Figure 5 - Identification of the stakeholders

9. Conclusion

The thesis aims to provide a novel approach at predicting clinical trial outcomes by integrating Large Language Models techniques and leveraging the Hierarchical Interaction Network. The goal of this research is to enhance the process of drug development, which is achieved by proposing the integration of LLMs predictive framework into the trial streamline. Our *HINTBasic* model's ability to handle a wide range of diverse datasets including integrate ICD codes, molecular data, and trial protocols resulted in significant enhancements in predictive performance. These advancements were exhibited by significant improvements in PR-AUC and ROC-AUC scores in comparison to baseline models, conveying robust

differentiation between successful and failed trials. Moreover, by comparing *HINTBasic* with *HINTPlus*, which includes enrollment, we saw improvements underscoring the importance of this new feature in achieving better predictions values, causing it to be our best model. By accelerating trial processes, the model is forecasted to decrease the duration of drug development by 15-30 %, resulting in savings of 6-12 months in earlier phases (Phase I and II) and a 20-25 % duration reduction in Phase III. Given its predictive power, our *HINTBasic* model can handle inefficiencies and significantly lower the high costs associated with failed trials. Indeed, the integration of the framework into a clinical workflow is estimated to reduce trials failures by 10 % to 15 %, potentially resulting in cost-savings equal to \$ 260 to \$ 390 million for each drug development process, if considered a total expenditure of \$ 2.6 billion. For instance, patient recruitment, which is a major cost center accounting for 30-40 % of the trial budget, could feature a cost reduction equal to 15 % due to a more targeted approach.

Although the model exhibits impactful benefits, it still must face challenges such as the model's reliance on high-quality. Data scarcity, specifically for rare diseases, is a significant risk to the prediction's performance.

Several stakeholders face significant implications, while also having to evaluate ethical considerations. While Researcher's and Health Care Provider's work is influenced by better resource allocation, while also reducing the initiation of unlikely successful trials, Pharmaceutical Companies benefit from accelerated drug launches, enhancing market competitiveness and ROI. Patients health outcome might also be improved through early risk identification, which can have an impact on their whole lifespan. All these stakeholders are affected by regulators, which can enhance the implementation of AI in clinical trials by collaboratively developing regulatory frameworks.

In terms of future outlooks, we believe that actions that might be performed on our current work, to accomplish further improvement, include dataset expansion, as diversity is required

to achieve generalizability and scalability.

An example of future work to be implemented is the development of personalized models to address clinical research for specific diseases and in techniques refinement to ensure granularity and accuracy of predictions. Additionally, ethical bias must be addressed, which can be executed through the integration of equity-focused frameworks and guidelines. Conclusively, the thesis depicts benefits and challenges associated with the adoption of machine learning techniques into the field of clinical trials.

Specifically, our work showcases the potential predictive ability of the *HINT* model to revolutionize the trial process, opening the frontier to a more operationally efficient and cost-effective drug development. Its implementation could lead the way for a major adoption of AI-driven models in the healthcare industry, ultimately enhancing the patients journey by reducing the leap time between treatments development and go-to market, while ensuring optimized efficiency.

10. Incorporating RAG Techniques for Predictive Modeling of Clinical Trial Outcomes

10.1 Introduction

Chopra et al. (2024) give reasonable arguments for why predictive models, and especially neural network models, face challenging implementation processes. The study credited the lack of explainability as one factor for why this is the case. In order to enhance the possibility implementation, this study focuses on the development of a Retrieval-Augmentation-Generation system which is supposed to contextualize predictions of the developed models by providing more information based on the user's needs, which has the potential to enhance confidence and understanding of the results.

Lewis et al. (2020) were early to introduce the Retrieval-Augmentation-Generation (RAG) architecture, which has been defined as a hybrid model architecture, combining pre-trained language models and retrieval-based external memory components for natural language generation. The architecture's key components are the following: The retrieval mechanism reflects the process of fetching information or documents from an external knowledge base. The foundation is a dense vector representation of the input query. The augmentation represents the retrieved information from the input query. The generator processes the augmented input and is often built using a transformer-based architecture in order to create responses based on the retrieved knowledge (Gupta et al. 2024).

Historically, a surge in development of traditional Natural Language Processing models has been observed. Nonetheless, these models greatly rely on training data, causing limitation when they are tasked to provide accurate answers regarding data the models were not trained on. As a result, traditional models could generate plausible yet incorrect information, reflecting the phenomenon of hallucinating large language models (Gupta et al. 2024).

In comparison to generative models, RAG has the ability to leverage external knowledge dynamically without requiring retraining, making it suitable for domains in which information is constantly evolving such as the medical field. However, this also requires it to handle growing datasets. While the retrieval mechanism is a powerful step, it can also lead to flawed results, which can be traced back to the vector representation of data, causing the retrieval of irrelevant documents. Furthermore, the generative module has the potential to be unable to incorporate retrieved information into its responses, which may create discrepancies between retrieval and generation. (Binns 2018). The objective is to incorporate RAG techniques in order to contextualize predictions leveraged by the *HINTPlus* and *HINTPlus per Disease*. Two RAG-Systems were developed in order to achieve this. The performance of the RAG-Systems was evaluated in comparison to each other. Moreover, it was tested how both systems react to general and more detailed prompts. Lastly, a Chatbot was employed to enable users to ask follow-up questions. The Generation system of the RAG-Systems alone as well as the Chatbot-enhanced RAG-Systems were evaluated in order to see if the system hallucinates answers.

10.2 Literature Review

Lu et al. (2024) introduced ClinicalRAG, a multi-agent pipeline in which every agent was in charge of a defined process stage, aiming to increase diagnosis accuracy. Not only was ClinicalRAG able to extract medical data from the user's input, it also had the capability to dynamically incorporate related medical knowledge in the text generation process, leading to the collection of high-quality information. Lu et al. (2024) observed that ClinicalRAG was able to significantly outperform traditional methods such as prompt learning, leading to enhanced reliability in clinical decision-making. However, limitations such as interpretability were stated, as they require attentions when it comes to an implementation.

Similar to Lu et al. (2024), Ong et al.'s (2024) employed a more general look on decision

support. Ong et al.'s (2024) study aimed at evaluating the effectiveness of RAG-based clinical decision support system regarding the identification of medication-related errors, by developing a simple and an advanced RAG-LLM. These two systems differed in numerous ways such as the chosen vector database, the retrieval scope or the embedding model. Their performances was evaluated on their own and as an assistive tool of a pharmacist. The study classified the simple RAG-LLM as the optimal choice, as it balanced accuracy, recall and precision. The usage as an assistive tool yielded superior results in comparison to the standalone variation, highlighting the potential of hybrid human-AI systems in clinical trial decision-making. However, the study concluded that data quality and scalability issues should be taken into account. While Ong et al. (2024) focused on patient safety in clinical care, Unlu et al. (2024) assessed the ability of a RAG-System to automate patient's eligibility screening by using electronic health records. The goal was the enhancement of accuracy, efficiency and cost-effectiveness regarding participant screening in comparison to traditional methods. They used GPT-4 for the generation aspect, while also comparing its performance to GPT-3.5 and manual screening done by staff. The results showed that their GPT-4 enhanced RAG-System demonstrated the highest accuracy, leading to relatively reduced screening costs. Although the GPT-4 enhanced RAG-System provided benefits such as accuracy and cost-effectiveness, it also came with challenges such as generalization and regulatory compliance.

Wang et al. (2023) took the implementation of GPT-based RAG systems one step further by exploring its integration with clinical-data into AI chatbots and comparing it to a baseline Chatbot, aiming to enhance patient guidance. AI chatbots such as ChatGPT demonstrated a deficiency of specificity in medical advice, which was conveyed by the frequent generation of generic responses. In contrast, the RAG-enhanced alternative showed enhanced accuracy and specificity by integrating relevant medical knowledge into its generated answers, entailing more detailed diagnostics and treatment recommendations. Yet, potential accelerated

treatments might also involve technical, ethical and scalability limitations.

10.3 Methodology

10.3.1 Database

With the goal of building a knowledge base, data curation steps were taken which were carried out for both RAG-Systems in the same manner. Contrary to conventional RAG-Systems, the knowledge bases are static. This was attributed to the fact that prediction as well as training data of both leveraged models were already up-to-date.

The derived results from both models, the *HINTPlus* and the *HINTPlus per Disease*, included NCT numbers and their respective predictions for all three phases separately. Since the results were saved in pkl.files, they were then converted to CSV files in order to make data handling easier. In addition, the developed benchmark dataset was utilized to retrieve the training data. Consequently, all columns which were not used for training the model were dropped, resulting in a filtered dataset containing 103,847 rows and 6 columns titled “NCT Number”, “Enrollment”, “Eligibility Criteria”, “Success/ Failure”, “SMILES” and “icd_code”. The CSV files containing NCT numbers and predictions were joined into one dataset. Since there are instances in which NCT numbers appeared in multiple phases, duplicated NCT numbers were dropped, ensuring the existence of unique NCT numbers and their respective predictions. Upon the described preparation steps, a left merge was performed. A significant amount of rows with no corresponding NCT numbers had missing value in the “Prediction” column, which were then eliminated. This resulted in the curation of a dataset with 12,817 rows for the *HINTPlus* RAG-System and 615 rows for the *HINTPlus per Disease* RAG-System while containing 7 columns respectively, serving as knowledge bases for both systems.


10.3.2 RAG Architectural Structure

The proposed RAG Architecture was built on insights from current literature such as Ong et al. (2024). As their simple RAG achieved good performance metrics, it will be the foundation

for the following proposed RAG Architecture. The curation of the dataset was followed by a data preprocessing step, in which all the columns of each row were batched into one single text string, containing all the information of each row. This simplified the system's performance of semantic similarity searches based on the users query, as every consolidated text serves as a representation of each trial. The clear data structure results in an efficient document retrieval (Burdick et al. 2021).

Once the preprocessing was completed, the sentence transformer “all-MiniLM-L6-v2” was employed, a well text encoder, ideal for detecting textual similarity. The model converted the combined text into numerical embeddings. Consequently, semantic relationships and contextual meanings were captured, enabling efficient document comparisons (Wang et al. 2020).

The numerical embeddings were then stored as vectors in a FAISS index, serving as an efficient library for similarity search. Its ability to retrieve similar documents based on the



embeddings by performing fast nearest neighbor searches was highly important, especially because the L2 distance metrics was also employed to handle the large data set. The same sentence transformer generated an embedding for a made query which prompted a search in the FAISS index, retrieving the top k, in this case 5, most relevant documents (Douze et al. 2024).

These documents were then transferred to OpenAI’s language model GPT-4o, with the ability to process complex information while generating clear and human like answers (Shahriar et al. 2024). The model was provided with the user’s query as well as the retrieved relevant documents in order to ensure a grounded answer. This was also reinforced by instructing the model to specifically use details from the retrieved documents. In order to enhance the user experience of the RAG-systems, a chatbot was added for both systems, allowing users to ask follow-up questions whilst maintaining context. Its foundation is a loop, handling the user input. It also reflected the RAG framework’s ability to retrieve the most relevant documents, to then generate a context-based response with GPT-4-o and outputs the answer (Wang et al. 2023). The same architectural structure was utilized for both RAG-Systems, the *HINTPlus* RAG-System and the *HINTPlus per Disease*. The following pipeline describes the proposed architecture:

Figure 7 - RAG-System Architecture

10.3.3 Experimental Setup and Hypotheses

The goal was to compare the ability of both developed RAG-Systems. First, we wanted to evaluate, if a disease-specific RAG-System depicts higher performance. Moreover, we wanted to test how the RAG-Systems reacted to Query-Engineering. The following hypotheses were tested to assess the retrieval component of both RAG-Systems: (1) The *HINTPlus per Disease* RAG-System will perform better than the *HINTPlus* RAG-System. (2) Queries detailing the user's role and goal of the query will yield a better performing RAG-System than a general query. In order to test these hypotheses, two queries were formulated. The first query entailed that the user wanted to know how high the enrollment number had to be regarding a Crohn's disease clinical trial in order to achieve a prediction of 0.8 or higher. The second query further included that the user is a medical researcher and that the answer would assist him in assessing success in future clinical trials. Apart from the added information about the user's role and the goal, the query stayed the same. Additionally, the Chatbot's potential to hallucinate answers was tested, by asking questions, which it could not be able to answer, based on the provided knowledge bases. In case of hallucination, we were refining the Chatbot by applying a feedback loop and thus giving users the ability to rate the Chatbots answers or refining the prompt and system message (Petridis et al. 2023).

10.4 Results

10.4.1 Retrieval

In order to test the hypotheses, we are following Ong et al.'s (2024) example by utilizing the same metrics. The first metrics is Precision@k, which quantifies the proportion of relevant documents within the top k retrieved documents. This metrics is especially useful to get insights into whether the top results presented are actually relevant. The second measure is Recall@k, measuring the proportion of relevant documents retrieved out of all relevant documents in the dataset. The third used metrics is the F1@k score, which is calculated as the

weighted average between Precision@k and Recall@k. Lastly, the Accuracy@k metric measures the proportion of correctly classified documents within the top k results. The following results were found by running the experiment of both RAG-Systems:

	<i>Performance Metrics Retrieval</i>			
	<i>RAG-System HINTPlus</i>		<i>RAG-System HINTPlus per Disease</i>	
	General Query	Specific Query	General Query	Specific Query
Precision@5	0.4000	0.6000	0.8000	0.8000
Recall@5	0.0104	0.0156	0.1600	0.1600
F1@5	0.0202	0.0304	0.2670	0.2670
Accuracy@5	0.4000	0.6000	0.8000	0.8000

Table 8 - Metrics measuring Retrieval Performance

While both RAG-Systems demonstrate relatively good performance in terms of Precision@5 and Accuracy@5, the same cannot be said about the Recall@5. This can be explained by the large number relevant data in the dataset, lying at 195. However, due to the set $k=5$, not all of them can be considered in the retrieval. A higher k could potentially lead to an improvement of this metric. Consequently, the low value in Recall@5 has also influenced the F1@5 score negatively. Overall, the resulting metrics suggested that the *HINTPlus per Disease* RAG-System demonstrated better performance overall, validating the first hypothesis. Additionally, queries detailing the user's role and goal led to higher performance achieved by the *HINTPlus* RAG-System. However, the same cannot be stated about the *HINTPlus per Disease* RAG-System, which demonstrated the same results for both the general and the detailed query. Consequently, the hypothesis can only be partially supported by the results of the *HINTPlus* RAG-System. The difference in performance might explainable by the complexity of the datasets used for both RAG-Systems. While the *HINTPlus* RAG-System contained a range of clinical trials regarding different diseases, the *HINTPlus per Disease* RAG-System dataset is merely focused on diseases about the human's digestive system.

10.4.2 Generation

We applied human evaluation in order to see if both RAG-Systems are generating correct outputs, while also answering the formulated query. Although, the RAG-Systems were at times prone to retrieving irrelevant documents, they consistently generated the correct outputs, responding to the query. In cases, in which irrelevant documents were retrieved, it was still able to filter them out in the generation step, leading to a correct output, which consistently answered the questions. This is especially important when considering challenges such as hallucination. Inputting queries, which were impossible to answer with only the documents inside the RAG-Systems, led to text generation which highlighted the fact that such questions cannot be answered based on the documents the RAG-Systems has been provided with. Overall, the generation demonstrated good performance consistently, while also highlighting its limitations thus avoiding hallucinations. However, the same cannot be said about the RAG-enabled Chatbot, as it has shown signs of hallucinated answers. Although, the same queries were prompted to the RAG-enhanced Chatbot, it was hallucinating answers which were not part of the provided datasets. Nonetheless, they were correct after verifying the information with *ClinicalTrial.gov* data. After refining the model's explicit system message and the prompt by highlighting the fact that the Chatbot is only supposed to use the retrieved documents while also applying a feedback loop which rewards the Chatbots highly ranked answers, no hallucination was detected after asking the same question again. Similar, to the RAG-System did the Chatbot inform the user about its incapability of answering the question due to insufficient data providing this information.

10.5 Discussion

Our results demonstrated that the retrieval mechanism of the *HINTPlus per Disease* RAG-System performed better than the *HINTPlus* RAG-System. Moreover, detailed queries including the user's role and goal only yielded a better performing retrieving for the

HINTPlus RAG-System. While the generating mechanism of both RAG-Systems showed great performance and no hallucination, this does not account for the RAG-enhanced-Chatbots. Both Chatbots showed signs of hallucination in their generated answers. However, refinements such led to no hallucinated answers after asking the same question again. It is important that these results have to be interpreted in the context of the following limitations.

10.5.1 Architectural Limitations

A higher setting of k could potentially lead to a higher Recall@ k score, which would in return influence the F1@ k score and thus the overall evaluation of both RAG-Systems. Additionally, the retrieval component could also benefit from a domain-specific embedding model, such as ClinicalBERT, as they demonstrate a higher understanding of the used vocabulary and the relationships within the data due to them being trained on it (Huang et al. 2019). Consequently, this could potentially impact the generational aspect positively as more relevant information is retrieved. The usage of FAISS for efficient similarity search could have also potentially influenced the findings. An alternative such as the HNSW (Hierarchical Navigable Small Word) algorithm, could be beneficial due to its ability to handle large datasets fast and efficiently (Malkov & Yashunin 2018). Another point to note is the lack of chunking in the proposed architecture. Chunking would provide the system with the capability of breaking down text elements or documents into smaller pieces before preprocessing them in the pipeline. This would provide the RAG-System to put more emphasis on specific information in large text and document, which could potentially increase the chance that the system ranks relevant information higher (Yepes et al. 2024) However, literature showed that a straightforward approach yielded better performance metrics which is why chunking was not applied in this case (Ong et al. 2024). The proposed RAG-Architecture lacks an important component of conventional RAG-Systems, which is the ability to

dynamically update its database. Since this work is an extension of set prediction derived from machine learning models, dynamic changes were considered.

10.5.2 Stakeholder Implementation

Although these constraints may impact the RAG-Systems performance and interpretability, there are also important implications. Machine learning predictive models are often seen as a “black box”, as they lack interpretability (Chopra et al. 2024). The proposed RAG-Systems could potentially enhance the confidence in the predictions made by the *HINTPlus* and *HINTPlus per Disease* models, due to its ability to answer a variety of queries, while also providing a contextualized answer, giving insights into the reasons of the prediction. The lack of hallucinated answers also supports this argument, especially since the RAG-Systems communicated its limitations when answering queries, going beyond the scope of the retrieved documents (Venkit et al. 2024). Enhanced usability by adding a Chatbot should also be considered, since this could potentially contribute to a higher understanding of the predictions, as users can ask follow up questions while still being able to see the whole conversation. The added feedback loop in order to reward the Chatbot contributes to the efforts of continuously improving the system. All of these points could help Healthcare Providers and Researchers to make informed decisions in their clinical trials, in terms of setting hypotheses, planning enrollment numbers comparing eligibility criteria which could also decrease errors per study. As a result, patients could benefit from receiving appropriate and safe care more efficiently. Companies would in return also benefit from the employment of the RAG-Systems as they could potentially influence drug discovery and thus the company’s ROI. The lack of hallucinated information ensures that the system is generating trustworthy information, which could potentially make Regulators gain trust in the system. Building on these limitations and implications, future research can explore several areas more in depth. Considering that this architecture does not link a database, which changes

dynamically, future research could design a RAG-architecture which hold its data up-to-date. Additionally, different embedding models, language models or vector databases could be employed with the goal to improve the system's performance.

10.6 Conclusion

The aim of this work was to build a RAG-System in order to leverage predictions made by the *HINTPlus* and *HINTPlus per Disease* model. Additionally, we created two Chatbots which were connected to their respective RAG-Systems. The two RAG-System's retrieval performance was compared overall and in dependance of the created query. Findings showed that the *HINTPlus per Disease* RAG-System performed better in terms of the retrieval component. Detailed queries only led to improved retrieval performance in the *HINTPlus per Disease* RAG-System. Generated answers were free of hallucinations. However, this changed after employing a RAG-enhanced Chatbot. Refinements by including a feedback mechanism and reformulating explicit service message and prompt caused the disappearance of hallucinated answers. Although the proposed RAG architecture could be improved in terms of its retrieving capability, it also provides beneficial implications for involved stakeholders in healthcare, such as the increased interpretability of predictions due to contextualized answers. Future research could refine this proposed architecture by implementing a key advantage of RAG-System, which is the ability to keep datasets up to date.

Bibliography

“BridgeBio Catapults After Snagging Approval for Its Rival to Pfizer’s Heart Drug - Bing.”

n.d. Bing. [https://www.bing.com/search?](https://www.bing.com/search?pglt=299&q=BridgeBio+Catapults+After+Snagging+Approval+For+Its+Rival+To+Pfizer%27s+Heart+Drug&cvid=7099336a5ec8472389f9c78d44c09d64&gs_lcrp=EgRlZGdlKgYIABBFgDkyBggAEEUYOdIBBzQ0MmowajGoAgCwAgA&FORM=ANNTA1&ucpdpc=UCPD&PC=ASTS)

[pglt=299&q=BridgeBio+Catapults+After+Snagging+Approval+For+Its+Rival+To+Pfizer%27s+Heart+Drug&cvid=7099336a5ec8472389f9c78d44c09d64&gs_lcrp=EgRlZGdlKgYIABBFgDkyBggAEEUYOdIBBzQ0MmowajGoAgCwAgA&FORM=ANNTA1&ucpdpc=UCPD&PC=ASTS](https://www.bing.com/search?pglt=299&q=BridgeBio+Catapults+After+Snagging+Approval+For+Its+Rival+To+Pfizer%27s+Heart+Drug&cvid=7099336a5ec8472389f9c78d44c09d64&gs_lcrp=EgRlZGdlKgYIABBFgDkyBggAEEUYOdIBBzQ0MmowajGoAgCwAgA&FORM=ANNTA1&ucpdpc=UCPD&PC=ASTS).

“Cost of Clinical Trials: A Breakdown (Infographic) – CRIO.” 2018. July 1, 2018. Accessed November 30th, 2024, <https://clinicalresearch.io/blog/running-a-study/cost-of-clinical-trials-breakdown/>.

“Understanding The Benefits Of Clinical Trials For Cancer.” 2021. Johns Hopkins Medicine. October 16, 2021. <https://www.hopkinsmedicine.org/health/treatment-tests-and-therapies/understanding-the-benefits-of-clinical-trials-for-cancer>.

Aliper, Alex, Roman Kudrin, Daniil Polykovskiy, Petrina Kamya, Elena Tutubalina, Shan Chen, Feng Ren, and Alex Zhavoronkov. 2023. “Prediction Of Clinical Trials Outcomes Based On Target Choice And Clinical Trial Design With Multi-Modal Artificial Intelligence.” *Clinical Pharmacology & Therapeutics* 114 (5): 972–80. <https://doi.org/10.1002/cpt.3008>.

Anagnostopoulos, Chris, David Champagne, Alex Devereson, Thomas Devenyns, and Heikki Tarkkila. 2023. “How Artificial Intelligence Can Power Clinical Development.” McKinsey & Company. November 22, 2023. Accessed November 30th 2024 <https://www.mckinsey.com/industries/life-sciences/our-insights/how-artificial-intelligence-can-power-clinical-development>.

Arrowsmith, John. 2011. “Phase III and submission failures: 2007–2010” *Nature Reviews Drug Discovery*, 10(2): 87–87. <https://www.nature.com/articles/nrd3375>.

Augustine, Erika F., Heather R. Adams, and Jonathan W. Mink. 2013. “Clinical Trials in Rare

Disease.” *Journal of Child Neurology* 28 (9): 1142–50.

<https://doi.org/10.1177/0883073813495959>.

Bieganek, Cameron, Constantin Aliferis, and Sisi Ma. 2022. “Prediction Of Clinical Trial Enrollment Rates.” *PLoS ONE* 17 (2): e0263193.

<https://doi.org/10.1371/journal.pone.0263193>.

Cassava Sciences, Inc. 2024. “Cassava Sciences Topline Phase 3 Data Did Not Meet Co-Primary Endpoints.” *GlobeNewswire News Room*, November 25, 2024.

<https://www.globenewswire.com/news-release/2024/11/25/2986578/8339/en/Cassava-Sciences-Topline-Phase-3-Data-Did-Not-Meet-Co-Primary-Endpoints.html>.

Chen, Tianqi & Guestrin, Carlos (2016). “XGBoost : A Scalable Tree Boosting System”.

Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794. <https://doi.org/10.1145/2939672.2939785>.

Chopra, Hitesh, None Annu, Dong K. Shin, Kavita Munjal, None Priyanka, Kuldeep Dhama, and Talha B. Emran. 2023. “Revolutionizing Clinical Trials: The Role Of AI in Accelerating Medical Breakthroughs.” *International Journal of Surgery* 109 (12): 4211–20.

<https://doi.org/10.1097/js9.0000000000000705>.

Cleophas, Ton J.M., and Ed M. De Vogel. 1998. “Crossover Studies Are A Better Format For Comparing Equivalent Treatments Than Parallel-group Studies.” *Pharmacy World & Science* 20 (3): 113–17. <https://doi.org/10.1023/a:1008626002664>.

Clinicaltrials.gov. (n.d.). <https://clinicaltrials.gov/about-site/about-ctg>.

Council, Science. 2023. “Potential Benefits And Limitations Of mRNA Technology For Vaccine Research And Development For Infectious Diseases And Virus-induced Cancers.” December 13, 2023. <https://www.who.int/publications/i/item/9789240084551>.

DiMasi, Joseph A., Henry G. Grabowski, and Ronald W. Hansen. 2016. “Innovation in the Pharmaceutical Industry: New Estimates of R&D Costs.” *Journal of Health Economics*

47 (February): 20–33. <https://doi.org/10.1016/j.jhealeco.2016.01.012>.

Edwards, Philip James, Ian Roberts, Mike J Clarke, Carolyn DiGuseppi, Benjamin Woolf, and Chloe Perkins. 2023. “Methods to Increase Response to Postal and Electronic Questionnaires.” Cochrane Library 2023 (11).

<https://doi.org/10.1002/14651858.mr000008.pub5>.

Faizullabhoj, Mariam, and Gauri Wani. 2024. “Clinical Trials Market - by Phase (I, II, III, IV), Study Design (Interventional, Observational, Expanded Access), Service Type (Outsourcing, In-house), Therapeutic Area (Oncology, Dermatology, Neurology, Cardiology) – Global Forecast (2024 – 2032).” Global Market Insights Inc. Accessed November 30th 2024 <https://www.gminsights.com/industry-analysis/clinical-trials-market>.

Fawcett, T. (2005). An introduction to ROC analysis. Pattern Recognition Letters, 27(8), 861–874. <https://doi.org/10.1016/j.patrec.2005.10.010>.

Financial Times. 2024. "Europe Falls Behind China in Playing Host to Clinical Drug Trials." Accessed November 30th 2024 <https://www.ft.com/content/3d861acb-8e7d-4157-b845-81124254da8a>.

Fu, Tianfan, Kexin Huang, Cao Xiao, Lucas M. Glass, and Jimeng Sun. 2022. “HINT: Hierarchical Interaction Network for Clinical-trial-outcome Predictions.” Patterns 3 (4): 100445. <https://doi.org/10.1016/j.patter.2022.100445>.

Gayvert, Kaitlyn M., Neel S. Madhukar, and Olivier Elemento. 2016. “A Data-Driven Approach To Predicting Successes And Failures Of Clinical Trials.” Cell Chemical Biology 23 (10): 1294–1301. <https://doi.org/10.1016/j.chembiol.2016.07.023>.

Getz, Kenneth A., Stella Stergiopoulos, Mary Short, Leon Surgeon, Randy Krauss, Sybrand Pretorius, Julian Desmond, and Derek Dunn. 2016. “The Impact of Protocol Amendments on Clinical Trial Performance and Cost.” Therapeutic Innovation & Regulatory Science 50 (4): 436–41. <https://doi.org/10.1177/2168479016632271>.

- Ghim, Jong-Lyul, and Sangzin Ahn. 2023. "Transforming Clinical Trials: The Emerging Roles Of Large Language Models." *Translational and Clinical Pharmacology* 31 (3): 131. <https://doi.org/10.12793/tcp.2023.31.e16>.
- Huang, Jin, and Ling, C. X. 2005. "Using AUC and Accuracy in Evaluating Learning Algorithms." *IEEE Transactions on Knowledge and Data Engineering*, 17(3): 299–310. doi: 10.1109/TKDE.2005.50.
- Huang, K., Altosaar, J., and Ranganath, R. 2019. "ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission." *arXiv (Cornell University)*, January. <https://doi.org/10.48550/arxiv.1904.05342>.
- John Hopkins Medicine. n.d. "Understanding the Benefits of Clinical Trials for Cancer." Accessed November 30th 2024 <https://www.hopkinsmedicine.org/health/treatment-tests-and-therapies/understanding-the-benefits-of-clinical-trials-for-cancer>.
- Kappen, Teus H., Wilton A. Van Klei, Leo Van Wolfswinkel, Cor J. Kalkman, Yvonne Vergouwe, and Karel G. M. Moons. 2018. "Evaluating The Impact Of Prediction Models: Lessons Learned, Challenges, And Recommendations." *Diagnostic and Prognostic Research* 2 (1). <https://doi.org/10.1186/s41512-018-0033-6>.
- Kavalci, Ece, and Anthony Hartshorn. 2023. "Improving Clinical Trial Design Using Interpretable Machine Learning Based Prediction Of Early Trial Termination." *Scientific Reports* 13 (1). <https://doi.org/10.1038/s41598-023-27416-7>.
- Kim, Eungdo, Jaehoon Yang, Sungjin Park, and Kwangsoo Shin. 2023. "Factors Affecting Success of New Drug Clinical Trials." *Therapeutic Innovation & Regulatory Science* 57 (4): 737–50. <https://doi.org/10.1007/s43441-023-00509-1>.
- Laakso, Mikael, and Bo-Christer Björk. 2012. "Anatomy Of Open Access Publishing: A Study Of Longitudinal Development And Internal Structure." *BMC Medicine* 10 (1). <https://doi.org/10.1186/1741-7015-10-124>.

Lo, Andrew W., Kien Wei Siah, and Chi Heem Wong. 2018. "Machine-Learning Models for Predicting Drug Approvals and Clinical-Phase Transitions." SSRN Electronic Journal, Modi, Shanu, William Jacot, Toshinari Yamashita, Joohyuk Sohn, Maria Vidal, Eriko Tokunaga, Junji Tsurutani, et al. 2022. "Trastuzumab Deruxtecan in Previously Treated HER2-Low Advanced Breast Cancer." *New England Journal of Medicine* 387 (1): 9–20. <https://doi.org/10.1056/nejmoa2203690>.

Montgomery, Alan A, Tim J Peters, and Paul Little. 2003. "Design, Analysis And Presentation Of Factorial Randomised Controlled Trials." *BMC Medical Research Methodology* 3 (1). <https://doi.org/10.1186/1471-2288-3-26>.

Murali, V. 2021. "Data - Predicting Clinical Trial Outcomes Using Drug Bioactivities." IEEE DataPort. December 27, 2021. <https://ieee-dataport.org/documents/data-predicting-clinical-trial-outcomes-using-drug-bioactivities>.

Nestor, Bret A., Matthew B. A. McDermott, Geeticka Chauhan, Tristan Naumann, Michael C. Hughes, A. Goldenberg, and M. Ghassemi. 2018. "Rethinking Clinical Prediction: Why Machine Learning Must Consider Year of Care and Feature Aggregation." 2018. <https://www.semanticscholar.org/paper/Rethinking-clinical-prediction%3A-Why-machine-must-of-Nestor-McDermott/393b41ea791f5a70ee87b361dec1970745d1908c#:~:text=This%20work%20augments%20MIMIC%20with%20the%20year%20in,aggregates%20of%20raw%20features%20significantly%20mitigate%20future%20deterioration>.

Office of the Commissioner. 2018. "Step 3: Clinical Research." U.S. Food And Drug Administration. January 4, 2018. <https://www.fda.gov/patients/drug-development-process/step-3-clinical-research>.

Parsons, Nick R., Joydeep Basu, and Nigel Stallard. 2024. "Group Sequential Designs For Pragmatic Clinical Trials With Early Outcomes: Methods And Guidance For Planning And Implementation." *BMC Medical Research Methodology* 24 (1).

<https://doi.org/10.1186/s12874-024-02174-w>.

Pytrial. “Welcome to PyTrial documentation!”. Accessed October 11, 2024.

<https://pytrial.readthedocs.io/en/latests>.

Reinisch, Michael, Jianfeng He, Chenxi Liao, Sauleh Ahmad Siddiqui, and Bei Xiao. 2024.

“CTP-LLM: Clinical Trial Phase Transition Prediction Using Large Language Models.”

arXiv.Org. August 20, 2024. <https://arxiv.org/abs/2408.10995>.

Saito, Takaya, and Rehmsmeier, Marc 2015. "The Precision-Recall Plot Is More Informative Than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets." PLoS One, 10(3): e0118432. DOI: 10.1371/journal.pone.0118432.

Singh, Nathan, Noelle V. Frey, Stephan A. Grupp, and Shannon L Maude. 2016. “CAR T Cell Therapy in Acute Lymphoblastic Leukemia and Potential for Chronic Lymphocytic Leukemia.” Current Treatment Options in Oncology 17 (6). <https://doi.org/10.1007/s11864-016-0406-4>.

Sokolova, Marina, and Lapalme, Guy 2009. "A Systematic Analysis of Performance Measures for Classification Tasks." Information Processing & Management, 45(4): 427–437. DOI:10.1016/j.ipm.2009.03.002.

Su, Qianmin, Gaoyi Cheng, and Jihan Huang. 2023. “A Review of Research on Eligibility Criteria for Clinical Trials.” Clinical and Experimental Medicine 23 (6): 1867–79. <https://doi.org/10.1007/s10238-022-00975-1>.

Ayala, Orlando, and Patrice Bechard. 2024. “Reducing Hallucination in Structured Outputs via Retrieval-Augmented Generation.” arXiv, January, 228–38. <https://doi.org/10.18653/v1/2024.naacl-industry.19>.

Binns, Rdp. 2018. “Fairness in Machine Learning: Lessons From Political Philosophy.” Journal of Machine Learning Research, January. <https://ora.ox.ac.uk/objects/uuid:2ff2785b-b0d4-447a-8326-a1fcc4c80840>.

Burdick, Laura, Jonathan K. Kummerfeld, and Rada Mihalcea. 2021. “To Batch or Not to Batch? Comparing Batching and Curriculum Learning Strategies Across Tasks and Datasets.” *Mathematics* 9 (18): 2234. <https://doi.org/10.3390/math9182234>.

Douze, Matthijs, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. “The Faiss Library.” arXiv (Cornell University), January. <https://doi.org/10.48550/arxiv.2401.08281>.

Gupta, Shailja, Rajesh Ranjan, and Surya Narayan Singh. 2024. “A Comprehensive Survey of Retrieval-Augmented Generation (RAG): Evolution, Current Landscape and Future Directions.” arXiv (Cornell University), October. <https://doi.org/10.48550/arxiv.2410.12837>.

Ji, Ziwei, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2022. “Survey of Hallucination in Natural Language Generation.” arXiv (Cornell University), January. <https://doi.org/10.48550/arxiv.2202.03629>.

Lewis, Patrick S. H., Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, et al. 2020. “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.” *Neural Information Processing Systems* 33 (May): 9459–74. <https://proceedings.neurips.cc/paper/2020/file/6b493230205f780e1bc26945df7481e5-Paper.pdf>.

Malkov, Yu A., and D. A. Yashunin. 2018. “Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42 (4): 824–36. <https://doi.org/10.1109/tpami.2018.2889473>.

Ong, Jasmine Chiat Ling, Liyuan Jin, Kabilan Elangovan, Gilbert Yong San Lim, Daniel Yan Zheng Lim, Gerald Gui Ren Sng, Yuhe Ke, et al. 2024. “Development and Testing of a Novel Large Language Model-Based Clinical Decision Support Systems for Medication Safety in 12 Clinical Specialties.” arXiv (Cornell University), January.

<https://doi.org/10.48550/arxiv.2402.01741>.

Petridis, Savvas, Ben Wedin, James Wexler, Aaron Donsbach, Mahima Pushkarna, Nitesh Goyal, Carrie J. Cai, and Michael Terry. 2023. “ConstitutionMaker: Interactively Critiquing Large Language Models by Converting Feedback Into Principles.” arXiv (Cornell University), January. <https://doi.org/10.48550/arxiv.2310.15428>.

Shahriar, Sakib, Brady D. Lund, Nishith Reddy Mannuru, Muhammad Arbab Arshad, Kadhim Hayawi, Ravi Varma Kumar Bevara, Aashrith Mannuru, and Laiba Batool. 2024. “Putting GPT-4o to the Sword: A Comprehensive Evaluation of Language, Vision, Speech, and Multimodal Proficiency.” *Applied Sciences* 14 (17): 7782.

<https://doi.org/10.3390/app14177782>.

Unlu, Ozan, Jiyeon Shin, Charlotte J Mailly, Michael F Oates, Michela R Tucci, Matthew Varugheese, Kavishwar Waghlikar, et al. 2024. “Retrieval Augmented Generation Enabled Generative Pre-Trained Transformer 4 (GPT-4) Performance for Clinical Trial Screening.” medRxiv (Cold Spring Harbor Laboratory), February.

<https://doi.org/10.1101/2024.02.08.24302376>.

Venkit, Pranav Narayanan, Tatiana Chakravorti, Vipul Gupta, Heidi Biggs, Mukund Srinath, Koustava Goswami, Sarah Rajtmajer, and Shomir Wilson. 2024. “An Audit on the Perspectives and Challenges of Hallucinations in NLP.” *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, January, 6528–48.

<https://doi.org/10.18653/v1/2024.emnlp-main.375>.

Wang, Calvin, Joshua Ong, Chara Wang, Hannah Ong, Rebekah Cheng, and Dennis Ong. 2023. “Potential for GPT Technology to Optimize Future Clinical Decision-Making Using Retrieval-Augmented Generation.” *Annals of Biomedical Engineering* 52 (5): 1115–18.

<https://doi.org/10.1007/s10439-023-03327-6>.

Wang, Wenhui, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020.


“MiniLM: Deep Self-Attention Distillation for Task-Agnostic Compression of Pre-Trained Transformers.” Neural Information Processing Systems 33 (February): 5776–88.

<https://proceedings.neurips.cc/paper/2020/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.


Yepes, Antonio Jimeno, Yao You, Jan Milczek, Sebastian Laverde, and Leah Li. 2024.

“Financial Report Chunking for Effective Retrieval Augmented Generation.” arXiv (Cornell University), February. <https://doi.org/10.48550/arxiv.2402.05131>.

Appendix



Graphic 1 - The graph shows an increasing trend in clinical studies over the years, with relative completion date according to ClinicalTrials.gov.




Graphic 2 - These graphs show the right-skewed distribution of the clinical trial duration in terms of year.

Files Missing Values			
File	Total Rows	Missing 'Enrollment'	%
phase_III_test.csv	3571	0	0.00%
phase_III_train.csv	16662	301	1.81%
phase_III_valid.csv	3570	1	0.03%
phase_II_test.csv	5796	0	0.00%
phase_II_train.csv	27047	609	2.25%
phase_II_valid.csv	5796	1	0.02%
phase_I_test.csv	4642	0	0.00%
phase_I_train.csv	21660	274	1.27%
phase_I_valid.csv	4642	0	0.00%
Total	93386	1186	1.27%

Table 1 - Number of rows with missing values

Enrollment Missing Counts by Study Status		
File Name	Study Status	Missing Counts
phase_I_train.csv	COMPLETED	263
phase_I_train.csv	TERMINATED	11
phase_II_train.csv	COMPLETED	574
phase_II_train.csv	SUSPENDED	1
phase_II_train.csv	TERMINATED	34
phase_III_train.csv	COMPLETED	280
phase_III_train.csv	TERMINATED	20
phase_III_train.csv	WITHDRAWN	1

Table 2 - Enrollment missing counts over Study Status




Graphic 3 - Enrollment missing values over years.


Estimate of F1 Score of 3 imputed datasets from MI			
	Phase I	Phase II	Phase III
M1	0.9309 ± 0.0026	0.8661 ± 0.0036	0.8995 ± 0.0040
M2	0.9319 ± 0.0024	0.8659 ± 0.0043	0.8986 ± 0.0042
M3	0.9335 ± 0.0037	0.8628 ± 0.0038	0.8993 ± 0.0041
Estimate	0.9321 ± 0.0034	0.8649 ± 0.0045	0.8991 ± 0.0041

Table 3 - Estimate of the multiple


imputations of the three imputed datasets using Rubin's Rule




Graphic 4 - Attention matrix heatmap in Phase I.




Graphic 5 - Attention matrix heatmap of Phase II




Graphic 6 - Attention matrix heatmap in Phase III.



Graphic 7 - Phase I – Performance Metrics Comparison



Graphic 8 - Phase II – Performance Metrics Comparison



Graphic 9 - Phase III – Performance Metrics Comparison

